Producing New Bijections from Old

(to appear in *Advances in Mathematics*)

April 22, 1994

David Feldman
University of New Hampshire

James Propp
Massachusetts Institute of Technology

ABSTRACT

For any purely combinatorial construction that produces new finite sets from given ones, a bijection between two given sets naturally determines a bijection between the corresponding constructed sets. We investigate the possibility of going in reverse: defining a bijection between the original objects from a given bijection between the constructed objects, in effect "canceling" the construction. We present a precise formulation of this question and then give a concrete criterion for determining whether such cancellation is possible. If the criterion is satisfied, then a cancellation procedure exists, and indeed, for many of the constructions that we have studied a polynomial-time cancellation procedure has been found; on the other hand, when the criterion is not satisfied, no procedure, however complicated, can perform the cancellation. The celebrated involution principle of Garsia and Milne fits into our rubric once it is recognized as a thinly disguised form of canceling disjoint union with a fixed set, itself a well–known technique that fits into our theoretical framework. We show that the construction that forms the Cartesian product of a given set with some other, fixed set cannot in general be canceled, but that it can be canceled in the case where the fixed set carries a distinguished element. We also show that the construction that takes the mth Cartesian power of a given set can be canceled, but that the construction that forms the power set of a given set cannot.

Key words: Bijection, bijective proof, G-set, Cartesian power, power set

1. Introduction

Among the various means of establishing combinatorial identities, many combinatorialists especially prize "bijective proofs": direct demonstrations by means of manifest bijections. The present status of the science of bijective proof may be likened to that of ruler-and-compass constructions prior to Gauss and Wantzel. Even with many ingenious examples, "bijective proof" remains an informal notion. The lack of suitable formalization prevents exploration of the limitations of "bijective proof," and so no results we know of may be interpreted to say, in any sense, that there exist analytically verifiable combinatorial identities admitting no bijective proof. Indeed the situation may well be sensitive to how "bijective proof" is made precise, much as angles may be trisected with marked straightedges, but not with unmarked straightedges.

In this paper, we show that just as algebraic methods (as embodied in Galois theory) can be brought to bear on questions of constructibility of geometrical objects by use of geometrical tools, so too can methods from algebra be applied to questions of constructibility of combinatorial objects via combinatorial tools such as bijections.

We do not offer a specific formalization of the notion of "bijective proof" in this paper. Rather we offer, and explore, a distinction between *weak* bijective proofs and *strong* bijective proofs that should be relevant in virtually any particular formal context. Let us now explain this distinction. A combinatorial identity concerns two indexed families of finite sets, say $\{A_i\}_{i\in I}$ and $\{B_i\}_{i\in I}$, and asserts that for all i, $|A_i| = |B_i|$. For our purposes nothing will be lost if we speak as though i were fixed, so hereafter we drop the subscript. Now, let $\mathcal{Q}(\cdot)$ be an arbitrary construction that produces new finite sets from old such that the cardinality of $\mathcal{Q}S$ determines the cardinality of S for all S. Then a manifest bijection between $\mathcal{Q}A$ and $\mathcal{Q}B$ also proves that $|A| = |B|$; we consider this a *weak bijective proof* that $|A| = |B|$, inasmuch as the fact that the cardinality of $\mathcal{Q}S$ determines the cardinality of S is not itself proved bijectively. A *strong bijective proof* that $|A| = |B|$ must exhibit a manifest bijection between the sets A and B themselves.

As we shall see, for some constructions $\mathcal{Q}(\cdot)$ it is always possible to *define* a bijection $f : A \to B$ *just* in terms of any given bijection $F : \mathcal{Q}A \to \mathcal{Q}B$, and

not taking into account of the special nature of the elements of the sets A and B. *Provided that the definition of f from F is within one's formal means, one is then guaranteed that a weak bijective proof using $\mathcal{Q}(\cdot)$ determines a strong bijective proof.* For our purposes, to say that the definition of f from F is within one's "formal means" is to say that f can be computed from F, in a scenario that we will make precise for the sake of definiteness (though the theorems that we prove are not overly sensitive to the precise details of the model). When the nature of $\mathcal{Q}(\cdot)$ permits us to compute f from F, we say that such a $\mathcal{Q}(\cdot)$ can be effectively canceled. We give examples of constructions $\mathcal{Q}(\cdot)$ that can be effectively canceled, along with others that cannot.

More generally, suppose that $\mathcal{Q}(\cdot,\cdot)$ is a construction that takes two arguments and that C is some sort of structured set, a notion that we shall make precise later. Suppose a bijection $F : \mathcal{Q}(A,C) \to \mathcal{Q}(B,C)$ is given. We derive a criterion for deciding whether a bijection $f : A \to B$ can be defined solely in terms of F and the given structure on C.

This criterion is algebraic in nature, and in specific applications it devolves into questions about permutation representations of finite groups. When the criterion is satisfied, a cancellation principle is guaranteed to exist, and we may proceed to seek it out. Indeed, we have typically found that when the cancelability criterion is satisfied, a cancellation procedure exists whose running-time is bounded by a polynomial function of the sizes of the sets in question.

If the criterion is not satisfied, then there is no point in searching for a cancellation procedure. However, even such negative results have positive import of a certain kind for the working combinatorialist. For instance, knowing that the power set construction cannot generally be canceled suggests the possibility of looking for interesting bijections between 2^A and 2^B even in situations where no interesting bijection between A and B is available. In any case, questions about cancelability, when run through the mill of our general-purpose machinery, give rise to interesting questions about permutation representations of finite groups, and thus can be a resource for algebraists in search of novel questions to ponder.

Our methods come from group theory and are purely mathematical, as opposed to metamathematical. Nevertheless it may be fruitful to interpret our

results as statements about combinatorics in various toposes, though we do not pursue this.

Our work can be seen as following in a tradition initiated by Adriano Garsia and Stephen Milne [4] in their study of the involution principle and continued by Herbert Wilf [12] and Basil Gordon [5] in their work on the principle of inclusion and exclusion.

Also of some relevance is work of László Lovász and Robert Appleson on finite sets with relational structures (see [1] and the references cited therein). Their work differs from ours in that we regard two finite structures as being equivalent if they are definable from one another, whereas those authors require that the structures actually be isomorphic. Nevertheless, there are many points of similarity between their papers and our paper, in terms of both methods used and results obtained, and it would be good to understand better the relationship between the two theories.

Our paper is organized as follows. In Section 2 we give preliminaries concerning group actions and fix some notation. Next, in Section 3 we introduce a protocol to make definite what we mean by effective cancellation. Then we prove Theorem 1, the Basic Equivariance Criterion, to formalize the intuitive notion "effectively cancelable combinatorial construction" (Section 4), and we give several variations thereof (Section 5). Theorem 2 recasts the Basic Equivariance Criterion in a more useful and more elegant form (Section 6). Section 7 digresses to treat some issues related to a variant form of cancellation, for the sake of completeness. Section 8 shows that issues of "uniformity" of cancellation procedures, which are not addressed elsewhere in the article, can be given a precise a formulation if one modifies the framework so that the role played by actions of groups is taken over by actions of monoids. The five sections that follow apply Theorem 2 to obtain concrete results, some positive and some negative, some new and some old: a positive result for cancellation of disjoint union with a fixed set (old, but we observe how it encompasses the celebrated involution principle of Garsia and Milne, thereby fitting that technique into our rubric); a negative result for Cartesian products (old, though not previously stated in this form in the literature, at least to our knowledge); a positive result for Cartesian products with pointed sets (new); a positive result for Cartesian powers (new); and a generally negative result for power sets (new) which is clarified by the

solution of an interesting group-theoretical problem. Section 14 gives a slick proof of a very general equivariance criterion. Section 15 situates our work in a more abstract, category-theoretic setting. Finally, Section 16 treats the relevance of our paper to issues in the foundations of mathematics. Open problems are interspersed throughout the paper.

It is not necessary to understand the foundational material in the first half of the paper to peruse the various explicit examples in Sections 9 through 13; indeed, these sections contain the main fruits of our work, and the reader may wish to sample them before studying the roots, trunk, and branches. As a middle course, the reader might start with Sections 2, 3, 4, and 6 before proceeding to Sections 9 through 13.

Some of our subsidiary propositions have trivial proofs that may not be trivial for the reader to find unaided; we have included compact versions of these proofs. However, rather than try to follow these verifications, the reader might find it more helpful to draw the appropriate diagrams and chase arrows in the usual way, consulting our proofs only as a last resort.

2. Preliminaries

We denote the integers by $\mathbf{Z}$, the natural numbers (including 0) by $\mathbf{N}$. All other sets considered are assumed finite.

A G-set is a set with an action by a group G (on the left). (For general background on group actions on finite sets, see Chapter 1 of [6]. If H is a subgroup of G, the coset-space G/H is also a G-set, under the definition $g(g'H) = (gg')H$. A bijection $f : A \to B$ between G-sets is a G-set isomorphism if $gf = fg$ for all $g \in G$. A G-orbit is a transitive G-set. Every G-set is a disjoint union of G-orbits. The stabilizer, Stab x, of x in a G-set is the group $\{g \mid gx = x\}$. If H is a subgroup of Stab x, we say H stabilizes x.

Let $\mathcal{O}$ be a G-orbit. For $x \in \mathcal{O}$, the map from $\mathcal{O}$ to $G/(\text{Stab } x)$ sending gx to $g(\text{Stab } x)$ for all $g \in G$ is a well-defined G-set isomorphism. If $\mathcal{O}$ contains an element whose stabilizer is H, we say that $\mathcal{O}$ is a G-orbit of H-type. Since Stab $gx = g(\text{Stab } x)g^{-1}$ for all $g \in G$, the type of an orbit is determined precisely up to conjugacy. So we have a correspondence between isomorphism classes of G-orbits and conjugacy classes of subgroups of G.

Let $L(G)$ be the lattice of subgroups of G. For any G-set S, we now define maps $\sigma_S, \sigma_S^+, \tau_S$ from $L(G)$ to $\mathbf{N}$. Let $\sigma_S(H)$ be the number of elements of S whose stabilizer is H. Let $\sigma_S^+(H)$ be the number of elements of S whose stabilizer contains H, i.e., the number of elements of S that are fixed by H. Let $\tau_S(H)$ be the number of G-orbits of S of H-type. Each function is constant on each conjugacy class of subgroups of G. Every $\mathbf{N}$-valued function that is constant on conjugacy classes of subgroups is $\tau_S(H)$ for some G-set S.

Each of the functions $\sigma_S(H)$, $\sigma_S^+(H)$, $\tau_S(H)$ determines the others, as follows.

By definition

$$\sigma_S^+(H) = \sum_{K \supseteq H} \sigma_S(K) \ .$$

On the other hand, writing this as

$$\sigma_S(H) = \sigma_S^+(H) - \sum_{K \supset H} \sigma_S(K)$$

allows us to compute σ_S from σ_S^+ by recursion down the lattice of subgroups (i.e., by Möbius inversion; see [6]).

For x in a G-orbit $\mathcal{O}$, Stab gx = Stab x iff $g(\text{Stab } x)g^{-1} = \text{Stab } x$ or equivalently $g \in N_G(\text{Stab } x)$, the normalizer of Stab x in G. This shows that the number of elements in an orbit of H-type with stabilizer H is $|N_G(H)/H|$, and so

$$\sigma_S(H) = |N_G(H)/H| \, \tau_S(H) \ .$$

Since the function $\tau_S(H)$ is clearly a complete numerical invariant for the isomorphism type of the G-set S, the functions $\sigma_S(H)$, $\sigma_S^+(H)$ are as well. Each turns out to be useful in its own way.

We write **Bij** for the category whose objects are finite sets and whose morphisms are bijections and $\text{Bij}(A, B)$ for the set of bijections from A to B. The undecorated expression $\text{Hom}(A, B)$ stands for the set of all functions from sets A to B; if we mean a Hom-set in some other category we indicate such via a subscript. The symmetric group on the set S is written Sym S. For the basic vocabulary of category theory, and much more, see [9].

3. Effective Cancellation Procedures

To make progress, we must analyze the intuitive notions "purely combinatorial construction" and "effectively cancelable combinatorial construction" in search of appropriate formal notions. A "purely combinatorial construction" $\mathcal{Q}(\cdot)$ takes a given set A and produces $\mathcal{Q}A$, a set which generally carries some further structure as well. Note that we need not specify the category in which $\mathcal{Q}(\cdot)$ takes its values.

Bijections $A \to B$ induce bijections $\mathcal{Q}A \to \mathcal{Q}B$ because we are assuming that $\mathcal{Q}(\cdot)$ takes no account of the names of particular elements. In fact, $\mathcal{Q}(\cdot)$ determines (but is not identical to) an endofunctor Q of the category of finite sets and bijective maps, i.e. a species in the sense of A. Joyal (see [7]). The usefulness of this notational distinction will become evident in Section 8, where combinatorial constructions $\mathcal{Q}(\cdot)$ are incarnated as functors $\overline{Q}$ between different respective categories. Moreover, we view combinatorial construction as an informal notion; category theory provides one language for formalizing this, but we prefer not to prematurely commit ourselves to a particular formalization, since there may be aspects of the naive notion that are not captured by any particular abstraction. We will continue to write $\mathcal{Q}A$ for QA when we wish to emphasize any extra structure the set carries by dint of its construction (e.g., the natural involution on $\mathcal{Q}A = A \times A$).

Henceforth we will assume that the endofunctor Q is faithful.

We give a precise meaning to *effective cancellation* in terms of a protocol involving two parties, Programmer and Machine. (These names are meant to emphasize the relationship between the parties rather than their human or mechanistic qualities.)

Fix a combinatorial construction $\mathcal{Q}(\cdot)$ and two disjoint finite sets A and B of cardinality n.

Assume (for now) that Programmer and Machine have agreed on names for the sets A and B themselves (but not for the elements they contain). Assume that they have also agreed on a common method by which to take a given set of symbols representing the elements of a set X and represent (possibly non-uniquely) the elements of $\mathcal{Q}X$ as finite strings which consist of symbols from a fixed finite common alphabet together with the symbols

that represent the elements of X. We will hereafter simplify matters by thinking of the elements of X as being themselves the symbols that are used to represent those elements, since this is merely a philosophical error, and not a mathematical mistake.

At the beginning of the protocol Programmer knows only the cardinality n, but Machine has rosters of elements for sets A, B and a bijection $F : \mathcal{Q}A \to \mathcal{Q}B$ (given in terms of the strings described above). We assume Machine has unbounded computational resources. In particular, Machine is able to make arbitrary choices.

During the protocol, Programmer sends instructions to Machine and never receives feedback. Programmer knows the general way in which Machine will present the elements of $\mathcal{Q}A$ and $\mathcal{Q}B$ in terms of the elements of A and B, respectively, but lacking rosters for A and B Programmer can never refer to the the individual members of the sets A and B. We may now make a

Definition: An *effective cancellation procedure* is a fixed sequence of instructions, depending only on n, for Programmer to send to Machine which upon execution always culminates in Machine's constructing a bijection $f : A \to B$ which depends only on F, not on arbitrary choices made by Machine.

While we do not stipulate that an effective cancellation procedure must return f when given $F = Qf$, from any given effective cancellation procedure we easily obtain an effective cancellation procedure that does have this property: Programmer simply asks Machine to return f should $F = Qf$ (at most one such f exists because Q is faithful) and to follow the given procedure otherwise.

4. The Basic Equivariance Criterion

An effective cancellation procedure for $\mathcal{Q}(\cdot)$ determines a function

$$\mathcal{F}_{A,B} : \mathrm{Bij}(QA, QB) \to \mathrm{Bij}(A, B) \ .$$

Since Programmer cannot refer to the names of elements of A and B, the function $\mathcal{F}_{A,B}$ must be invariant under relabelings. More exactly, the group

Sym $A \times$ Sym B acts on $\mathrm{Bij}(A,B)$ by

$$(\rho_1,\rho_2)f = \rho_2 f \rho_1^{-1}$$

and on $\mathrm{Bij}(QA,QB)$ by

$$(\rho_1,\rho_2)F = \overline{\rho_2} F \overline{\rho_1^{-1}},$$

where $\overline{\rho}$ is shorthand for $Q\rho$. This will always be the action that we intend, even in the case where A and B intersect. (One might at first think of introducing a compatibility condition between ρ_1 and ρ_2, but this is inappropriate; to make sense of the concept of disjoint union, for instance, one has to conceive of there being two copies of any elements in the intersection.) Invariance under relabelings amounts to

$$\mathcal{F}_{A,B}((\rho_1,\rho_2)F) = \mathcal{F}_{A,B}(\overline{\rho_2} F \overline{\rho_1^{-1}}) = \rho_2 \mathcal{F}_{A,B}(F)\rho_1^{-1} = (\rho_1,\rho_2)\mathcal{F}_{A,B}(F)\ ,$$

that is, $\mathcal{F}_{A,B}$ must be (Sym $A \times$ Sym B)-equivariant.

The crucial but perhaps less obvious fact is the converse: the mere existence of a (Sym $A \times$ Sym B)-equivariant map in $\mathrm{Hom}(\mathrm{Bij}(QA,QB),\mathrm{Bij}(A,B))$ (the set of arbitrary set maps from $\mathrm{Bij}(QA,QB)$ to $\mathrm{Bij}(A,B)$) guarantees an effective cancellation procedure. Though we give a proof right now in the spirit of our development so far, in Section 14 below the reader may find a short, conceptual proof of a much more general result.

First a trivial but useful

Lemma 1 *Given $a \in \mathrm{Bij}(A,\tilde{A})$ and $b \in \mathrm{Bij}(B,\tilde{B})$, if $\mathcal{F}_{A,B}$ is* (Sym $A \times$ Sym B)*-equivariant, then $\mathcal{F}_{\tilde{A},\tilde{B}}$, given by*

$$\mathcal{F}_{\tilde{A},\tilde{B}}(F) = b\mathcal{F}_{A,B}(\overline{b^{-1}} F \overline{a})a^{-1}, \quad F \in \mathrm{Bij}(Q\tilde{A},Q\tilde{B})$$

is (Sym $\tilde{A} \times$ Sym $\tilde{B}$)*-equivariant.*

Proof:

$$\mathcal{F}_{\tilde{A},\tilde{B}}(\overline{\rho_2} F \overline{\rho_1^{-1}}) = b\mathcal{F}_{A,B}(\overline{b^{-1}\rho_2} F \overline{\rho_1^{-1} a})a^{-1} = b\mathcal{F}_{A,B}(\overline{(b^{-1}\rho_2 b)b^{-1}} F \overline{a(a^{-1}\rho_1^{-1}a)})a^{-1}$$

$$= b(b^{-1}\rho_2 b)(\mathcal{F}_{A,B}(\overline{b^{-1}} F \overline{a}))(a^{-1}\rho_1^{-1}a)a^{-1} = \rho_2 \mathcal{F}_{\tilde{A},\tilde{B}}(F)\rho_1^{-1}\ .\Box$$

Here is the effective cancellation procedure:

First, Programmer instructs Machine to well–order the set $\mathrm{Hom}(\mathrm{Bij}(Q\mathbf{n}, Q\mathbf{n}), \mathrm{Bij}(\mathbf{n}, \mathbf{n}))$, where $\mathbf{n} = \{1, \ldots, n\}$. In more detail: Machine uses the well–ordering on $\mathbf{n}$ to represent each element of $Q\mathbf{n}$ by a *unique* string, say the first lexicographically from among the shortest strings representing that element previously. Since each element of $Q\mathbf{n}$ is now represented by a *unique* string, the lexicographic ordering of these strings determines a well–ordering of $Q\mathbf{n}$. Machine is now able to represent each bijection in $\mathrm{Bij}(Q\mathbf{n}, Q\mathbf{n})$ by a unique string of elements from $Q\mathbf{n}$: the domain of such a bijection is well–ordered, so Machine can simply list the respective elements of the range that the elements of the domain are mapped to under the bijection. Moreover, Machine can well–order these strings lexicographically, using the well–ordering on the range. This well–orders $\mathrm{Bij}(Q\mathbf{n}, Q\mathbf{n})$. $\mathrm{Bij}(\mathbf{n}, \mathbf{n})$ and then $\mathrm{Hom}(\mathrm{Bij}(Q\mathbf{n}, Q\mathbf{n}), \mathrm{Bij}(\mathbf{n}, \mathbf{n}))$ are well–ordered similarly.

Next, Machine uses the well–ordering on $\mathcal{H} = \mathrm{Hom}(\mathrm{Bij}(Q\mathbf{n}, Q\mathbf{n}), \mathrm{Bij}(\mathbf{n}, \mathbf{n}))$ to pick the least map in $\mathcal{H}$ which is $(\mathrm{Sym}\ \mathbf{n} \times \mathrm{Sym}\ \mathbf{n})$-equivariant; call it $\mathcal{F}_{\mathbf{n},\mathbf{n}}$. Since we assume $\mathrm{Hom}(\mathrm{Bij}(QA, QB), \mathrm{Bij}(A, B))$ contains a $(\mathrm{Sym}\ A \times \mathrm{Sym}\ B)$-equivariant map we are assured that $\mathrm{Hom}(\mathrm{Bij}(Q\mathbf{n}, Q\mathbf{n}), \mathrm{Bij}(\mathbf{n}, \mathbf{n}))$ contains a $(\mathrm{Sym}\ \mathbf{n} \times \mathrm{Sym}\ \mathbf{n})$-equivariant map by Lemma 1.

Programmer now instructs Machine to choose (arbitrary!) bijections $a : \mathbf{n} \to A$ and $b : \mathbf{n} \to B$. Another application of Lemma 1 yields a $(\mathrm{Sym}\ A \times \mathrm{Sym}\ B)$-equivariant map, call it $\mathcal{F}_{a,b}$. Finally Machine outputs $\mathcal{F}_{a,b}(F)$.

The crucial point is that the choice of a and b is immaterial. If $\alpha : \mathbf{n} \to A$ and $\beta : \mathbf{n} \to B$ are different choices,

$$\mathcal{F}_{a,b}(F) = b\mathcal{F}_{\mathbf{n},\mathbf{n}}(\overline{b^{-1}}F\overline{a})a^{-1} = b\mathcal{F}_{\mathbf{n},\mathbf{n}}(\overline{(b^{-1}\beta)\beta^{-1}}F\overline{\alpha(\alpha^{-1}a)})a^{-1}$$

$$= b(b^{-1}\beta)(\mathcal{F}_{\mathbf{n},\mathbf{n}}(\overline{\beta^{-1}}F\overline{\alpha}))(\alpha^{-1}a)a^{-1} = \beta(\mathcal{F}_{\mathbf{n},\mathbf{n}}(\overline{\beta^{-1}}F\overline{\alpha}))\alpha^{-1} = \mathcal{F}_{\alpha,\beta}(F)\ .$$

We have proved the following

Proposition 1 *Bijections between $\mathcal{Q}(A)$ and $\mathcal{Q}(B)$ can be effectively canceled iff there exists a* $(\mathrm{Sym}\ A \times \mathrm{Sym}\ B)$*-equivariant map from* $\mathrm{Bij}(QA, QB)$ *to* $\mathrm{Bij}(A, B)$. □

Comment From the point of view of our motivation it is proper to speak of canceling constructions, not bijections. Nevertheless, when a construction can't be canceled as a whole, it is still possible that particular bijections F between $\mathcal{Q}(A)$ and $\mathcal{Q}(B)$ can be used to define particular bijections f between A and B. By a harmless abuse of language, we say that such an F can be effectively canceled.

Comment Note that if A' and B' have the same cardinality as A and B, then bijections between $\mathcal{Q}A'$ and $\mathcal{Q}B'$ can be effectively canceled if and only if bijections between $\mathcal{Q}A$ and $\mathcal{Q}B$ can be effectively canceled. That is, the existence or nonexistence of an effective cancellation procedure depends only on the cardinality n of the sets A and B. On the other hand, if A' and B' have a different cardinality than A and B, it can happen that cancellation is possible for one pair but not for the other.

Comment When a (Sym $A \times$ Sym B)-equivariant map from Bij(QA, QB) to Bij(A, B) exists, the proof of Proposition 1 gives an effective cancellation procedure which unfortunately requires Machine to spend time exponential in $|QA|$. However, as we have remarked earlier, in almost every case where we have first verified, using group-theoretic methods, that the Basic Equivariance Criterion is satisfied, a polynomial-time cancellation procedure has subsequently emerged.

Problem 1 *Do their exist cancelable constructions which nevertheless cannot be canceled in time polynomial in the cardinalities of the sets?*

Some combinatorial constructions make use of an auxiliary set C, e.g., $\mathcal{Q}A = A \times C$ (Cartesian product). Since Programmer has thus far been permitted free reference to the manner in which elements of $\mathcal{Q}X$ are specified in terms of elements of X, in particular Programmer has had free reference by name to the elements of any auxiliary set. We now seek a refinement of Proposition 1 where such reference is restricted.

We model constructions with auxiliary sets as bifunctors $Q(\cdot, \cdot)$. In particular, this implies that Sym A and Sym C both act on $Q(A, C)$, and that these actions *commute*.

For the time being let us assume that C is disjoint from A and B.

When Programmer is permitted no reference even to the elements of C, then any cancellation procedure must be invariant under relabeling the elements of C as well. More generally, we may wish to equip the set C with a "structure" $\mathcal{S}$ and then allow Programmer reference only to $\mathcal{S}$, but not reference by name to the elements of C (unless, of course, these names are an aspect of $\mathcal{S}$).

It turns out that our present purposes do not require us to define "structure" formally or to specify a mechanism by which Programmer will communicate to Machine about $\mathcal{S}$. All we need assume for now is that fixing such an $\mathcal{S}$ determines a subgroup Aut C of Sym C to be interpreted as the group of bijections which preserve $\mathcal{S}$.

The following result is fundamental to the rest of our work. One may prove it by mimicking the proof of Proposition 1; it is also an immediate corollary of the much more general Theorem 13 proved in Section 14.

Theorem 1 (Basic Equivariance Criterion) *Bijections between $Q(A,C)$ and $Q(B,C)$ may be effectively canceled without reference to the elements of the auxiliary set C iff there exists a map from* Bij$(Q(A,C),Q(B,C))$ *to* Bij(A,B) *which is* (Sym $A \times$ Sym $B \times$ Aut C)*-equivariant.* □

5. Other Equivariance Criteria

Theorem 13 in Section 14 below is substantially more general than Theorem 1 above. Here we describe some corollaries of Theorem 13 that play a supporting role in our main applications. Other corollaries can be found in Section 14. Like Theorem 1, each of these results may also be proved by mimicking the proof of Proposition 1.

Heretofore we have allowed Programmer to distinguish between the sets A and B. We will speak of *anonymous effective cancellation* when reference by name to the sets A and B themselves is forbidden. An anonymous effective cancellation procedure must be invariant not only under relabeling the elements of A and B separately, but also under the simultaneous relabeling of the elements of A as elements of B and vice-versa.

Let Sym$(A;B)$ be the subgroup of Sym $(A \dot{\cup} B)$ consisting of elements that

respect the partition. We regard elements of $\mathrm{Sym}(A;B)$ as pairs of bijections, either of the form $(\alpha : A \to A,\ \beta : B \to B)$ or $(\alpha : A \to B,\ \beta : B \to A)$. Thus $\mathrm{Sym}(A;B)$ is a two-fold extension of $\mathrm{Sym}\ A \times \mathrm{Sym}\ B$. We extend the $(\mathrm{Sym}\ A \times \mathrm{Sym}\ B \times \mathrm{Aut}\ C)$-action on $\mathrm{Bij}(Q(A,C), Q(B,C))$ to a $(\mathrm{Sym}(A;B) \times \mathrm{Aut}\ C)$-action by setting

$$(\alpha : A \to B,\ \beta : B \to A,\ \eta)F = \overline{\alpha}\eta F^{-1}\eta^{-1}\overline{\beta^{-1}}\ .$$

We obtain a $\mathrm{Sym}(A;B)$-action on $\mathrm{Bij}(A,B)$ similarly.

Corollary 1 (Equivariance Criterion for Anonymous Effective Cancellation) *Bijections between $Q(A,C)$ and $Q(B,C)$ may be effectively canceled without reference to A and B by name iff there exists a* $(\mathrm{Sym}(A;B) \times \mathrm{Aut}\ C)$*-equivariant map from* $\mathrm{Bij}(Q(A,C), Q(B,C))$ *to* $\mathrm{Bij}(A,B)$.

Let $\mathcal{Q}_1(\cdot,\cdot)$ and $\mathcal{Q}_2(\cdot,\cdot)$ be combinatorial constructions of two arguments. Constructions $\mathcal{Q}_1(\cdot,\cdot)$ and $\mathcal{Q}_2(\cdot,\cdot)$ are represented by bifunctors from the category of finite sets and bijections to itself, say $Q_1(\cdot,\cdot)$ and $Q_2(\cdot,\cdot)$. Suppose that $|A| = |B|$ iff $|Q_1(A,B)| = |Q_2(A,B)|$. The natural $\mathrm{Sym}\ A \times \mathrm{Sym}\ B$ action on $\mathrm{Bij}(Q_1(A,B), Q_2(A,B))$ and previous considerations lead to:

Corollary 2 (Equivariance Criterion for Bifunctors) *Bijections between $Q_1(A,B)$ and $Q_2(A,B)$ can be used to effectively define bijections between A and B iff there exists a* $(\mathrm{Sym}\ A \times \mathrm{Sym}\ B)$*-equivariant map from* $\mathrm{Bij}(Q_1(A,B), Q_2(A,B))$ *to* $\mathrm{Bij}(A,B)$.

For example, we might have

$$Q_1(A,B) = (A \times A)\dot{\cup}(B \times B) \text{ and } Q_2(A,B) = (A \times B)\dot{\cup}(B \times A).$$

An explicit procedure for constructing a bijection between A and B in this case is given in the middle of Section 11, after we have acquired more tools.

Suppose that effective cancellation is impossible for $\mathcal{Q}(\cdot,\cdot)$ applied to A, B and C with a given set of restrictions on Programmer. How might we measure the extent of the failure? One way is to ask *which* bijections on the constructed objects exhibit sufficient lack of symmetry that they allow one

to define bijections between the original objects. (We apply this method in our analysis of the power-set construction in Section 13.) A different approach is to look for ways to make Programmer's job easier by permitting Programmer to refer to some particular extra (individual or collective) structure $\mathcal{S}$ on the sets A, B and C.

Giving examples of potentially useful structures is easy. Here are two examples of individual structure, and two of collective structure.

1. A distinguished point in the set C.
2. Cyclic orderings on the sets A and B.
3. An orientation of the formal simplex with vertex-set $A \dot{\cup} B$.
4. A bijection between sets $A \dot{\cup} B$ and C.

(Of course, we have already considered the utility of reference by name to the elements of C and the utility of reference by name to the sets A and B, or put differently, a distinguished element in the set $\{A, B\}$.)

For an approach to formalizing the intuitive notion "structure," see Section 14. For now, let us assume only that a "structure" $\mathcal{S}$ on A, B and C determines, if nothing else, an automorphism group $\mathcal{G}_{\mathcal{S}} \subseteq (\mathrm{Sym}(A; B) \times \mathrm{Sym}\ C)$ whose elements are the maps which preserve $\mathcal{S}$. Furthermore, we assume having the ability to communicate about $\mathcal{S}$ allows Programmer to instruct Machine to construct $\mathcal{G}_{\mathcal{S}}$ (e.g., by checking each element of $(\mathrm{Sym}(A; B) \times \mathrm{Sym}\ C)$ for preservation of $\mathcal{S}$). An important special case is when $\mathcal{S}$ is a distinguished element in the set $\{A, B\}$ (which allows Programmer to refer A and B by name). Then $\mathcal{G}_{\mathcal{S}}$ is $\mathrm{Sym}\ A \times \mathrm{Sym}\ B \times \mathrm{Sym}\ C$.

Corollary 3 (Equivariance Criterion with Collective Structure) *Bijections between $Q(A, C)$ and $Q(B, C)$ may be effectively canceled making use of a structure $\mathcal{S}$ on A, B and C iff there exists a map from* $\mathrm{Bij}(Q(A, C), Q(B, C))$ *to* $\mathrm{Bij}(A, B)$ *which is $\mathcal{G}_{\mathcal{S}}$-equivariant.*

Note that the utility of a structure $\mathcal{S}$ for purposes of cancellation depends only on $\mathcal{G}_{\mathcal{S}}$. In particular, a bijection F in $\mathrm{Bij}(Q(A, C), Q(B, C))$ can be canceled with reference to the structure $\mathcal{S}$ if and only if the intersection of $\mathrm{Stab}\ F$ with $\mathcal{G}_{\mathcal{S}}$ (viewed as subgroups of $\mathrm{Sym}(A; B) \times SymC$) is contained in the stabilizer of some $\mathrm{Stab}\ f$.

We can think of the elements of $(\mathrm{Sym}(A;B) \times \mathrm{Sym}\ C)$ that preserve some element of $\mathcal{H} = \mathrm{Bij}(Q(A,C), Q(B,C))$, but preserve no element of $\mathrm{Bij}(A,B)$, as constituting *obstructions* to the construction of a bijection between A and B. The set of obstructions is a comprehensive measure of the difficulty of effectively canceling $\mathcal{Q}(\cdot,\cdot)$ applied to A, B and C. Alternatively, one can focus on the totality $\mathcal{T}_{\mathcal{Q}}$ of subgroups $\mathcal{G}$ of $(\mathrm{Sym}(A;B) \times \mathrm{Sym}\ C)$ that contain no obstructions, i.e., that have the property that $\mathcal{H}$ contains a $\mathcal{G}$-equivariant map. The set of obstructions determines the collection $\mathcal{T}_{\mathcal{Q}}$, and vice versa, but philosophically it may be helpful to focus on $\mathcal{T}_{\mathcal{Q}}$. In particular, the maximal subgroups in $\mathcal{T}_{\mathcal{Q}}$ correspond to the minimal amounts of extra structure that suffice to facilitate cancellation.

Similar measures may be developed relative to the presence of some given structure $\mathcal{S}$. When the availability of $\mathcal{S}$ does not suffice for effectively canceling $\mathcal{Q}(\cdot,\cdot)$ applied to A, B and C, define $\mathcal{T}_{\mathcal{Q}}(\mathcal{S})$ to be the set of subgroups $\mathcal{G}$ of $\mathcal{G}_{\mathcal{S}}$ such that $\mathcal{H}$ contains a $\mathcal{G}$-equivariant map. Then $\mathcal{T}_{\mathcal{Q}}(\mathcal{S})$ measures the amount of further structure needed in addition to $\mathcal{S}$ to effect cancellation.

Problem 2 *Develop a theory of the collections $\mathcal{T}_{\mathcal{Q}}(\mathcal{S})$ that can arise when cancellation fails. In particular, can they be explicitly described in the case when Q is the power set functor?*

Whenever A, B and C are not (as previously assumed) pairwise disjoint, obtain an equivalent cancellation problem as follows: make disjoint copies of the sets, but take as *extra structure* the relations between the disjoint sets induced by the identity relations between the original sets. Extra structure can only help us cancel, so the assumption that our original sets were disjoint cannot make cancellation easier. We make use of this later when we give counterexamples to the existence of effective cancellation procedures in which A, B and C are not pairwise disjoint. Such examples are often easier to think about.

6. Main Theorem

With Theorem 2 below, the question of the existence of an effective cancellation procedure for a construction $\mathcal{Q}(\cdot,\cdot)$ involving a structured auxiliary set C is reduced to group theory.

Lemma 2 *Let G be a (finite) group and S and T be G-sets. Then the following are equivalent:*

(i) There is a G-equivariant map h from S to T.

(ii) For every $s \in S$ there is a $t \in T$ such that Stab $s \subseteq$ Stab t.

Proof

(i)⇒(ii): Stab $s \subseteq$ Stab $h(s)$ since $gs = s$ implies $gh(s) = h(gs) = h(s)$.

(ii)⇒(i): Pick a representative $s_{\mathcal{O}}$ of each orbit $\mathcal{O}$ of S. Set $h(s_{\mathcal{O}}) = t_{\mathcal{O}}$ for some element $t_{\mathcal{O}}$ of T satisfying Stab $s_{\mathcal{O}} \subseteq$ Stab $t_{\mathcal{O}}$. Extend h to a map from S to T by sending $h(gs_{\mathcal{O}})$ to $gt_{\mathcal{O}}$. To see that h is well–defined, suppose $g_1 s_{\mathcal{O}} = g_2 s_{\mathcal{O}}$. Then $g_2^{-1} g_1 \in$ Stab $s_{\mathcal{O}} \subseteq$ Stab $t_{\mathcal{O}}$, so that $g_1 t_{\mathcal{O}} = g_2 t_{\mathcal{O}}$ as desired. To see that h is equivariant, note that for all $g' \in G$, $h(g'(gs_{\mathcal{O}})) = h((g'g)s_{\mathcal{O}}) = g'gt_{\mathcal{O}} = g'(h(gs_{\mathcal{O}}))$. □

Returning now to the Basic Equivariance Criterion, we now see that there is an equivariant map

$$\mathcal{F}_{A,B} : \mathrm{Bij}(Q(A,C), Q(B,C)) \to \mathrm{Bij}(A,B)$$

iff for every $F \in \mathrm{Bij}(Q(A,C), Q(B,C))$, Stab F is contained in one of the groups Stab f with $f \in \mathrm{Bij}(A,B)$.

In the spirit of the previous remark, let us say that a particular bijection $F : Q(A,C) \to Q(B,C)$ is *effectively cancelable* if Stab F is contained in one of the groups Stab f with $f \in \mathrm{Bij}(A,B)$. Thus the construction $Q(\cdot,\cdot)$ is effectively cancelable iff every $F : Q(A,C) \to Q(B,C)$ is effectively cancelable. Issues of "uniformity" only play a role if we make further restrictions such as resource bounds on Machine. See also Section 8 below.

Each of the sets A, B, $Q(A,C)$, $Q(B,C)$, $\mathrm{Bij}(A,B)$ and $\mathrm{Bij}(Q(A,C), Q(B,C))$ carries a (Stab F)-action obtained by restricting the natural (Sym $A \times$ Sym $B \times$ Sym C)-action. Specifically: for $a \in A$, $(\rho_1, \rho_2, \eta)(a) = \rho_1(a)$; for $b \in B$, $(\rho_1, \rho_2, \eta)(b) = \rho_2(b)$; for $\overline{a} \in Q(A,C)$, $(\rho_1, \rho_2, \eta)(\overline{a}) = \overline{\rho_1}\eta(\overline{a})$; for $\overline{b} \in Q(B,C)$ $(\rho_1, \rho_2, \eta)(\overline{b}) = \overline{\rho_2}\eta(\overline{b})$; for $f \in \mathrm{Bij}(A,B)$, $(\rho_1, \rho_2, \eta)(f) = \rho_2 f \rho_1^{-1}$; and for $F \in \mathrm{Bij}(Q(A,C), Q(B,C))$, $(\rho_1, \rho_2, \eta)F = \eta\overline{\rho_2}F\overline{\rho_1^{-1}}\eta^{-1}$.

In particular we have

Lemma 3 *The following two statements are equivalent:*

(i) The bijection $f : A \to B$ satisfies Stab $F \subseteq$ *Stab* f.

(ii)The bijection $f : A \to B$ is an isomorphism of (Stab F)*-sets.*

Proof Statement (i) says $f = \rho_2 f \rho_1^{-1}$ for all $(\rho_1, \rho_2, \eta) \in$ Stab F; statement (ii) says $f\rho_1 = \rho_2 f$ for all $(\rho_1, \rho_2, \eta) \in$ Stab F. □

Lemma 4 *Any bijection $F : Q(A,C) \to Q(B,C)$ is itself an isomorphism of* (Stab F)*-sets.*

Proof Since $\eta^{-1}\overline{\rho_2^{-1}}F\overline{\rho_1}\eta = (\rho_1^{-1}, \rho_2^{-1}, \eta^{-1})F = F$,

$$F((\rho_1, \rho_2, \eta)(\overline{a})) = F\overline{\rho_1}\eta(\overline{a}) = \overline{\rho_2}\eta(F(\overline{a})) = (\rho_1, \rho_2, \eta)(F(\overline{a}))\ .\Box$$

Given a group G and G-sets A and C, $Q(A,C)$ is a G-set in a natural way, since Q is a bifunctor. Suppose C carries a structure $\mathcal{S}$. Recall that Aut C is the subgroup of Sym C consisting of bijections which preserve $\mathcal{S}$. A G-action on C is *$\mathcal{S}$-compatible* if it is induced by a homomorphism $G \to$ Aut C.

We can now state (and prove) the main result of this section, which is used in nearly everything that follows.

Theorem 2 *Suppose the finite set C carries a structure $\mathcal{S}$. Then the following are equivalent:*

(I) For all finite groups G, all finite G-sets A and B, and for all $\mathcal{S}$-compatible G-actions on C, $Q(A,C)$ and $Q(B,C)$ are isomorphic if and only if A and B are isomorphic;

(II) The construction $\mathcal{Q}(\cdot,\cdot)$ can be effectively canceled using reference to the structure $\mathcal{S}$ on C.

Proof (I)⇒(II):
Let $G =$ Sym $A \times$ Sym $B \times$ Aut C and fix a bijection $F : Q(A,C) \to Q(B,C)$. By Lemma 4, F is an isomorphism of (Stab F)-sets. As such, (I) guarantees the existence of some (Stab F)-isomorphism $f : A \to B$. Then Lemma 3 guarantees that Stab $F \subseteq$ Stab f. Now, letting F vary, we see that (ii) of Lemma 2 is satisfied, so there exists a (Sym $A \times$ Sym $B \times$ Aut C)-equivariant

map from $\mathrm{Bij}(Q(A,C),Q(B,C))$ to $\mathrm{Bij}(A,B)$. By the Basic Equivariance Criterion, such a map guarantees us an effective cancellation procedure for Q using reference to the structure $\mathcal{S}$ on C.

(II)$\Rightarrow$(I):
Suppose (I) fails. Then we may fix a finite group G, *non*-isomorphic G-sets A and B and a homomorphism $G \to \mathrm{Aut}\ C$ such such that $Q(A,C)$ and $Q(B,C)$ are isomorphic G-sets. Fix a G-set isomorphism $F : Q(A,C) \to Q(B,C)$. Naturally the image of G in $\mathrm{Sym}\ A \times \mathrm{Sym}\ B \times \mathrm{Aut}\ C$ is a subgroup of $\mathrm{Stab}\ F$. Since A and B are non-isomorphic as G-sets, they are *a fortiori* non-isomorphic as ($\mathrm{Stab}\ F$)-sets. Since no map $f : A \to B$ is a ($\mathrm{Stab}\ F$)-isomorphism, Lemma 3 guarantees that there can be no f such that $\mathrm{Stab}\ F \subseteq \mathrm{Stab}\ f$. Lemma 2 thus precludes the existence of a ($\mathrm{Sym}\ A \times \mathrm{Sym}\ B \times \mathrm{Aut}\ C$)-equivariant map from $\mathrm{Bij}(Q(A,C),Q(B,C))$ to $\mathrm{Bij}(A,B)$, and so there can be no effective cancellation procedure, even using the structure $\mathcal{S}$ on C. □

7. Anonymous Cancellation

Theorem 2 gives a group-theoretic condition for effective cancellation when reference by name to the sets A and B is permitted. A further group-theoretic condition must be met if the construction $\mathcal{Q}(\cdot,\cdot)$ is to be *anonymously* effectively canceled using reference to the structure $\mathcal{S}$ on C.

Recall the definition of anonymous effective cancellation from Section 5. Recall also that we say that a G-action on C is *$\mathcal{S}$-compatible* if it is induced by a homomorphism $G \to \mathrm{Aut}\ C$.

Theorem 3 *Assume* $\mathcal{Q}(\cdot,\cdot)$ can *be effectively canceled using structure* $\mathcal{S}$ *on* C. *Then the following are equivalent:*

(I) For all finite groups G, finite G-sets A, $\mathcal{S}$-compatible G-actions on C, and G-set automorphisms F of $Q(A,C)$, and for all $(\alpha,\beta,g) \in \mathrm{Sym}\ A \times \mathrm{Sym}\ A \times G$, if $\overline{\alpha}g = F\overline{\beta}gF$ then there exists a G-set automorphism f of A such that $\alpha g = f\beta g f$.

(II) The construction $\mathcal{Q}(\cdot,\cdot)$ can be anonymously *effectively canceled with reference to the structure $\mathcal{S}$ on C.*

Note that the bijections α and β are *not* presumed to be G-set isomorphisms.

Proof (I)⇒(II):
Suppose (II) fails. By the Equivariance Criterion for Anonymous Effective Cancellation, there must be a bijection $F : Q(A,C) \to Q(B,C)$ such that Stab F ($\subseteq$ (Sym$(A;B) \times$ Aut C)) is not contained in Stab f for any $f : A \to B$. Now set $G =$ Stab $F \cap$ (Sym $A \times$ Sym $B \times$ Aut C) and fix an element $(\alpha : A \to B, \beta : B \to A, g) \in$ (Stab F) $\setminus G$. Since (α, β, g) fixes F, $\overline{\alpha} g F^{-1} g^{-1} \overline{\beta^{-1}} = F$, that is, $\overline{\alpha} g = F \overline{\beta} g F$. But (α, β, g) generates Stab F over G, so (α, β, g) cannot fix any $f : A \to B$ stabilized by G. That is, if f is a G-set isomorphism, then $\alpha g f^{-1} g^{-1} \beta^{-1} \neq f$ or equivalently $\alpha g \neq f \beta g f$. By the previous Theorem, together with our assumption that $\mathcal{Q}(\cdot,\cdot)$ can be effectively canceled, G does fix some $f_0 : A \to B$. Identifying A and B along f_0 contradicts (I).

(II)⇒(I):
Suppose (I) fails. Then there is a finite group G, a finite G-set A, an $\mathcal{S}$-compatible G-action on C, a G-set automorphism F of $Q(A,C)$ and an element $(\alpha, \beta, g) \in$ Sym $A \times$ Sym $A \times G$ such that $\overline{\alpha} g = F \overline{\beta} g F$, but for no G-set automorphism f of A do we have $\alpha g = f \beta g f$. Let B be a G-set isomorphic to A and $i : A \to B$ be a G-set isomorphism. (It is helpful to think of A and B as being two copies of the "same" set, and i as the map that identifies the two copies.) Then the bijection $\overline{i} F : Q(A,C) \to Q(B,C)$ is stabilized by $(i\alpha, \beta i^{-1}, g) \in$ Sym$(A;B) \times$ Aut C, but no G-set isomorphism $if : A \to B$ (they are all of this form) is stabilized by $(i\alpha, \beta i^{-1}, g)$. Thus Stab $\overline{i} F \not\supseteq$ Stab if, so by Lemma 2 there can be no (Sym$(A;B) \times$ Aut C)-equivariant map, and no anonymous effective cancellation procedure. □

Problem 3 *Find methods for checking the group-theoretic condition for anonymous effective cancellation.*

8. Uniform Cancellation and Monoids

This section develops criteria for a more stringent notion of cancellation germane only to a special class of combinatorial constructions. Since much of this section parallels earlier passages, we will include fewer details than before.

As we have seen, a cancellation for each *individual* bijection $F : QA \to QB$ implies the existence of an effective cancellation for the *entire* construction $\mathcal{Q}$, one algorithm which Machine may apply uniformly to cancel all such bijections, regardless of the cardinalities of A and B. Uniformity may become a non-trivial issue in a variety of ways, for example, by imposing resource bounds on Machine's algorithm. Nevertheless, we shall henceforth use the term *uniform* to refer specifically to a a certain type of structural compatibility between cancellations involving different cardinalities. This notion of uniformity will be *prima facie* independent of algorithmic concerns.

As we consider this new notion of *uniform cancellation*, parallels and contrasts with our previous development will emerge. Monoids and categories of M-sets will play the role formerly played by groups and categories of G-sets. The analogue of Theorem 2 for uniform cancellation does not guarantee effectivity.

We begin by defining *uniform combinatorial constructions*. For simplicity, we restrict ourselves here to combinatorial constructions $\mathcal{Q}(\cdot)$ that take just a single argument. Throughout the discussion, A, A', B and B' will be finite sets. Observe that since each element of $\mathcal{Q}(A)$ is an equivalence class of strings of elements from A and a fixed alphabet, an equivalence relation on A induces one on $\mathcal{Q}(A)$.

Let us say that a combinatorial construction $\mathcal{Q}(\cdot)$ is *uniform* if it satisfies the conditions

> (i) If $A \subset A'$, then $\mathcal{Q}(A) \subset \mathcal{Q}(A')$;
> (ii) If A is a set of representatives for an equivalence relation on A'', then $\mathcal{Q}(A)$ is a set of representatives for the induced equivalence relation on $\mathcal{Q}(A'')$.

A uniform combinatorial construction $\mathcal{Q}(\cdot)$ determines an endofunctor $\overline{Q}$: **FinSet** $\to$ **FinSet**, where **FinSet** is the category of finite sets and arbitrary set maps. This endofunctor $\overline{Q}$ extends the endofunctor Q : **Bij** $\to$ **Bij** that one always has.

Examples The n^{th} Cartesian power construction, which associates to a set A the set of strings of length n of elements from A, determines a functor

$\overline{Q} : \mathbf{FinSet} \to \mathbf{FinSet}$, but the construction $\mathrm{LinOrd}(\cdot)$, which associates to a set A the set of linear orders of A, does not satisfy (i).

Next we explain the notions of *effective uniform cancellation procedure* and then *uniform cancellation* for a uniform combinatorial construction. Fix a bijection $F : \overline{Q}A \to \overline{Q}B$. Suppose that Machine is actually given, not F, but a bijection $F' : \overline{Q}A' \to \overline{Q}B'$, such that F' is compatible with F in the sense that there exist maps $a : A' \to A$ and $b : B' \to B$ such that $(\overline{Q}b)F' = F(\overline{Q}a)$. Note that neither Machine nor Programmer is given the maps a and b, or even the sets A and B. Programmer's goal is now to direct Machine in the construction of a bijection $f' : A' \to B'$ which is to be likewise compatible with some bijection $f : A \to B$ in the sense that $bf' = fa$. Moreover we impose a natural *uniformity requirement*: f should depend only on F, not on the choice of F', a or b. If this is possible for all A, B and F we say that $\mathcal{Q}(\cdot)$ admits an *effective uniform cancellation procedure*.

We indicate our interpretation of the uniformity requirement. If the maps a and b above are surjections, we may regard the sets A' and B' as consisting of multiple names for the elements of A and B, the bijection $F' : \overline{Q}A' \to \overline{Q}B'$ as a redundant representation of the bijection $F : \overline{Q}A \to \overline{Q}B$, and the bijection $f' : A' \to B'$ as a correspondingly redundant representation of a bijection $f : A \to B$. If the maps a and b are injections, one may regard the sets A and B as padded versions of the sets A' and B', the bijection $F' : \overline{Q}A \to \overline{Q}B$ as a representation of the bijection $F : \overline{Q}A \to \overline{Q}B$ padded with superfluous data, and the bijection $f' : A \to B$ as a correspondingly padded representation of a bijection $f : A' \to B'$. Thus cancellation must respect not just relabeling, but collapsing and padding as well.

An effective uniform cancellation procedure for $\mathcal{Q}(\cdot)$ determines, for any pair of finite sets A and B, maps

$$\mathcal{F}_{A,B} : \mathrm{Bij}(\overline{Q}A, \overline{Q}B) \to \mathrm{Bij}(A, B)$$

satisfying an analogue of equivariance:

> (*) For all $a : A \to A$ and $b : B \to B$, $b\mathcal{F}_{A,B}(F) = \mathcal{F}_{A,B}(F)a$ whenever $(\overline{Q}b)F = F(\overline{Q}a)$.

More generally, the uniformity requirement entails that the family of maps satisfy the following compatibility condition:

(**) Given maps $a : A' \to A$, $b : B' \to B$ and bijections $F' : \overline{Q}A' \to \overline{Q}B'$, $F : \overline{Q}A \to \overline{Q}B$ such that $(\overline{Q}b)F' = F(\overline{Q}a)$, then $b\mathcal{F}_{A',B'}(F') = \mathcal{F}_{A,B}(F)a$.

Let us call such a compatible family of maps a *uniform cancellation* for $\overline{Q}$. We emphasize that a uniform cancellation need not be computable.

We now introduce a new formalism particularly suited to the study of uniform cancellation.

To the functor $\overline{Q} : \mathbf{FinSet} \to \mathbf{FinSet}$ we associate the comma category $(\overline{Q} \downarrow \overline{Q})$. We follow the terminology and notation of [9], but our treatment is self-contained. An object in $(\overline{Q} \downarrow \overline{Q})$ is a triple $\langle A, B, F \rangle$, where A and B are objects of **FinSet** and F is a function $F : \overline{Q}A \to \overline{Q}B$. A morphism in $(\overline{Q} \downarrow \overline{Q})$ from $\langle A', B', F' \rangle$ to $\langle A, B, F \rangle$ is a pair of **FinSet** morphisms $\langle a, b \rangle = \langle a : A' \to A,\ b : B' \to B \rangle$ such that $F(\overline{Q}a) = (\overline{Q}b)F'$.

Morphisms may be visualized as commutative squares:

$$\begin{array}{ccc} \overline{Q}A' & \xrightarrow{F'} & \overline{Q}B' \\ \overline{Q}a \downarrow & & \downarrow \overline{Q}b \\ \overline{Q}A & \xrightarrow{F} & \overline{Q}B \end{array} .$$

We denote by $\mathbf{FinSet}_{\overline{Q}}$ the full subcategory of $(\overline{Q} \downarrow \overline{Q})$ consisting of objects $\langle A, B, F \rangle$ where F happens to be a bijection.

The forgetful functor from $(\overline{Q} \downarrow \overline{Q})$ (or $\mathbf{FinSet}_{\overline{Q}}$) to $\mathbf{FinSet} \times \mathbf{FinSet}$ takes objects $\langle A, B, F \rangle$ to pairs $\langle A, B \rangle$, and morphisms $\langle a, b \rangle$ to $\langle a, b \rangle$. In particular, the notation $\langle a, b \rangle$ does not by itself determine a morphism in $(\overline{Q} \downarrow \overline{Q})$ or $\mathbf{FinSet}_{\overline{Q}}$ since it specifies neither F' nor F. This ambiguity turns out to be very convenient.

In the case of the identity functor $\mathrm{Id} : \mathbf{FinSet} \to \mathbf{FinSet}$, we obtain categories $(\mathrm{Id} \downarrow \mathrm{Id})$ and $\mathbf{FinSet}_{\mathrm{Id}}$. In this setting a *uniform cancellation* is merely a functor $\mathcal{F} : \mathbf{FinSet}_{\overline{Q}} \to \mathbf{FinSet}_{\mathrm{Id}}$ which commutes with the forgetful functors just described, or using the ambiguity just noted, a functor which satisfies $\mathcal{F}\langle a, b \rangle = \langle a, b \rangle$.

If $\mathbf{C}$ is a subcategory of $\mathbf{FinSet}_{\overline{Q}}$, we call a functor $\mathcal{F} : \mathbf{C} \to \mathbf{FinSet}_{\mathrm{Id}}$ satisfying $\mathcal{F}\langle a, b\rangle = \langle a, b\rangle$ a $\mathbf{C}$-uniform cancellation. If $\mathbf{C}$ happens to be the full category of $\mathbf{FinSet}_{\overline{Q}}$ determined by a class of objects C, we also refer to this sort of functor as a C-uniform cancellation.

Observe that a uniform cancellation is determined by how it maps objects of $\mathbf{FinSet}_{\overline{Q}}$ to objects of $\mathbf{FinSet}_{\mathrm{Id}}$. Also, for a mapping $\mathcal{F}$ from objects of $\mathbf{FinSet}_{\overline{Q}}$ to objects of $\mathbf{FinSet}_{\mathrm{Id}}$ to give a uniform cancellation, it suffices for $\mathcal{F}$ to be well-behaved on morphisms in the following sense: given a morphism $\langle a, b\rangle : F' \to F$ in $\mathbf{FinSet}_{\overline{Q}}$, $\langle a, b\rangle : \mathcal{F}F' \to \mathcal{F}F$ is a morphism in $\mathbf{FinSet}_{\mathrm{Id}}$. Functoriality is then automatic because morphisms compose according to their names, and these are preserved.

(We caution the reader that the commutative squares above may also be viewed as morphisms in a different category, this time with the vertical arrows the objects and the diagrams composed side by side. We most definitely *don't* ask for $\mathcal{F}$ to be functorial in this sense, but see Section 15.)

Note that the objects in $\mathbf{FinSet}_{\overline{Q}}$ (resp. $\mathbf{FinSet}_{\mathrm{Id}}$) of the form $\langle A, B, F\rangle$ (resp. $\langle A, B, f\rangle$) are just the elements of $\mathrm{Bij}(\overline{Q}A, \overline{Q}B)$ (resp. $\mathrm{Bij}(A, B)$). Moreover, a uniform cancellation $\mathcal{F}$ restricted to $\mathrm{Bij}(\overline{Q}A, \overline{Q}B)$ is the familiar map $\mathcal{F}_{A,B}$.

As before, we write $\mathbf{n}$ for $\{1, \ldots, n\}$. Let $\mathbf{O}$ (resp. $\mathbf{O_j}$) be the full subcategory of $\mathbf{FinSet}_{\overline{Q}}$ determined by the set $\bigcup_{\mathbf{n=1}}^{\infty} \mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$ (resp. $\bigcup_{\mathbf{n=1}}^{\mathbf{j}} \mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$). Every morphism in $\mathbf{FinSet}_{\overline{Q}}$ is isomorphic (as an object in the category whose objects are $\mathbf{FinSet}_{\overline{Q}}$-morphisms and whose morphisms are commuting squares) to a morphism in $\mathbf{O}$. Thus uniform cancellation $\mathcal{F} : \mathbf{FinSet}_{\overline{Q}} \to \mathbf{FinSet}_{\mathrm{Id}}$ exists if and only an $\mathbf{O}$-uniform cancellation exists.

For each j, suppose $\mathcal{F}_j$ is an $\mathbf{O_j}$-uniform cancellation. When $j > n$, $\mathcal{F}_j\langle \mathbf{n}, \mathbf{n}, F\rangle$ is an element of the *finite* set $\mathrm{Bij}(\mathbf{n}, \mathbf{n})$. This shows by compactness that some subsequence of the $\mathcal{F}_j$ converges, in a suitable sense, to an $\mathbf{O}$-uniform cancellation $\mathcal{F}$. Since the finite subcategories of $\mathbf{O}$ are precisely the subcategories of the $\mathbf{O_j}$'s, we conclude that an $\mathbf{O}$-uniform cancellation exists if and only if $\mathbf{o}$-uniform cancellations exist for every finite subcategory $\mathbf{o}$ of $\mathbf{O}$. This also follows from the compactness theorem of first-order logic.

Next we introduce a condition on the functor $\overline{Q}$ from which it will follow that

a uniform cancellation exists if and only if $\mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$-uniform cancellations exist for all $\mathbf{n}$.

Let us say that a $\mathbf{FinSet}_{\overline{Q}}$-object $\langle A', B', F' \rangle$ *covers* object $\langle A, B, F \rangle$ if there are morphisms

$$\langle A, B, F \rangle \xrightarrow{\langle i_A, i_B \rangle} \langle A', B', F' \rangle \xrightarrow{\langle j_A, j_B \rangle} \langle A, B, F \rangle$$

with composition the identity map on $\langle A, B, F \rangle$.

Let us say that the functor $\overline{Q}$ is *levelable* if for every finite subcategory $\mathbf{o}$ of $\mathbf{FinSet}_{\overline{Q}}$, there exists $\mathbf{n}$ such that each object of $\mathbf{o}$ is covered by some object in $\mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$.

Lemma 5 *Suppose $\overline{Q}$ is a levelable functor with $\mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$-uniform cancellations for all $\mathbf{n}$. Then $\overline{Q}$ has a uniform cancellation.*

Proof It is enough to see that $\overline{Q}$ has an $\mathbf{o}$-uniform cancellation whenever $\mathbf{o}$ is a finite subcategory of $\mathbf{O}$. Suppose each object $\langle \mathbf{m}, \mathbf{m}, F_m \rangle$ in $\mathbf{o}$ may be covered by a corresponding object $\langle \mathbf{n}, \mathbf{n}, \tilde{F}_m \rangle$ in $\mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$. To each object $\langle \mathbf{m}, \mathbf{m}, F_m \rangle$ in $\mathbf{o}$ we must associate an object $\langle \mathbf{m}, \mathbf{m}, f_m \rangle$ in $\mathbf{FinSet}_{\mathrm{Id}}$ in such a way that if

$$\langle \mathbf{m}, \mathbf{m}, F_m \rangle \xrightarrow{\langle a, b \rangle} \langle \mathbf{p}, \mathbf{p}, F_p \rangle$$

is a morphism in $\mathbf{o}$, then

$$\langle \mathbf{m}, \mathbf{m}, f_m \rangle \xrightarrow{\langle a, b \rangle} \langle \mathbf{p}, \mathbf{p}, f_p \rangle$$

is a morphism in $\mathbf{FinSet}_{\mathrm{Id}}$.

We assemble the given data into the following diagram, where $\overline{a} = i_2 a j_3$ and $\overline{b} = i_4 b j_1$:

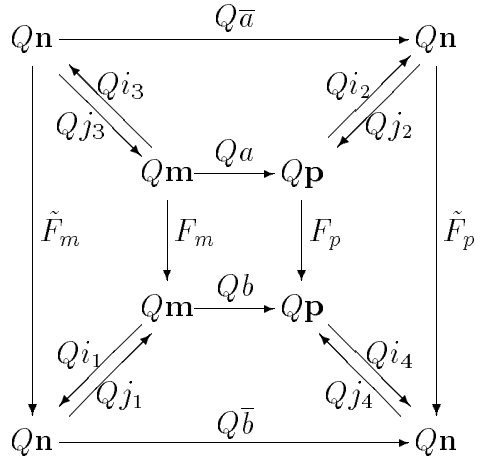

The outer square commutes because each of the inner squares does. Thus

$$\langle \mathbf{n}, \mathbf{n}, \tilde{F}_m \rangle \xrightarrow{\langle a,b \rangle} \langle \mathbf{n}, \mathbf{n}, \tilde{F}_p \rangle$$

is a morphism in $\mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$. By assumption there is a $\mathrm{Bij}(\overline{Q}\mathbf{n}, \overline{Q}\mathbf{n})$-uniform cancellation, so the outer square of the following diagram commutes:

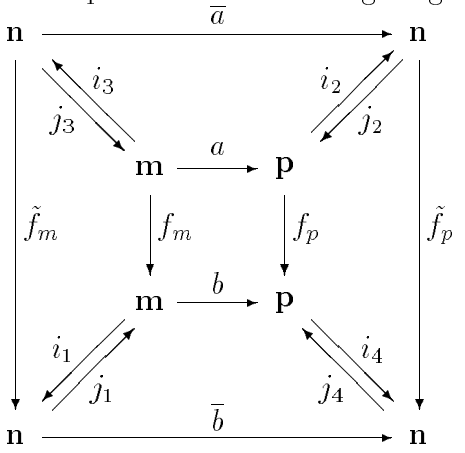

Moreover, on setting $f_m = j_1 \tilde{f}_m i_3$ and $f_p = j_4 \tilde{f}_p i_2$ we see that since $\overline{b}\tilde{f}_m = \tilde{f}_p \overline{a}$, we have $j_4 \overline{b} \tilde{f}_m i_3 = j_4 \tilde{f}_p \overline{a} i_3$, $j_4 \overline{b} i_1 j_1 \tilde{f}_m i_3 = j_4 \tilde{f}_p i_2 j_2 \overline{a} i_3$, $b j_1 \tilde{f}_m i_3 = j_4 \tilde{f}_p i_2 a$ and $b f_m = f_p a$. □

Examples Let us call a function $j : A' \to A$ *flat* if $|j^{-1}(a)|$ is constant as a varies over A. Any functor $\overline{Q} : \mathbf{FinSet} \to \mathbf{FinSet}$ that takes flat functions to flat functions is levelable. To see this, suppose $\{F_k : Q\mathbf{m_k} \to Q\mathbf{m_k}\}$ is the set of morphisms of a finite subcategory $\mathbf{o}$ of $\mathbf{FinSet}_{\overline{Q}}$. Let $\mathbf{n}$ be a common

multiple of the the $\mathbf{m_k}$. As long as $j_k : \mathbf{n} \to \mathbf{m_k}$ is a flat function that is a left inverse to an injection $i_k : \mathbf{m_k} \to \mathbf{n}$, it is easy to find a cover

$$\langle \mathbf{m_k}, \mathbf{m_k}, F_k \rangle \overset{\langle i_A, i_B \rangle}{\longrightarrow} \langle \mathbf{n}, \mathbf{n}, \tilde{F}_k \rangle \overset{\langle j_k, j_k \rangle}{\longrightarrow} \langle \mathbf{m_k}, \mathbf{m_k}, F_k \rangle \ .$$

In particular, any construction of the form $\mathcal{Q}(A) = \dot{\cup}_{\ell=1}^{n} C_\ell A^\ell$ is uniform and levelable.

Now let Func A be the monoid of arbitrary self-maps of A. Even though we have no action now of Func $A \times$ Func B on Bij$(\overline{Q}A, \overline{Q}B)$, given $F \in$ Bij$(\overline{Q}A, \overline{Q}B)$, we will still write Stab F for the monoid consisting of all pairs $(a : A \to A,\ b : B \to B)$ such that $(\overline{Q}b)F = F(\overline{Q}a)$. Similarly, for $f \in$ Bij(A, B), we will now write Stab f for the monoid consisting of all pairs (a, b) such that $bf = fa$. Given a bijection $F : A \to B$, as before, A and B are naturally Stab F sets. In this context it is immediate that F itself is a Stab F-isomorphism. It is also immediate that a Bij$(\overline{Q}B, \overline{Q}A)$-uniform cancellation $\mathcal{F}_{A,B}$ exists just in case each Stab F is included in some Stab f.

The rest of the proof of the following theorem follows the lines of the proof of Theorem 2:

Theorem 4 *Let $\mathcal{Q}(\cdot)$ be a uniform, levelable construction. Then the following are equivalent:*

(I) For all finite monoids M, all finite M-sets A and B, $\overline{Q}(A)$ and $\overline{Q}(B)$ are isomorphic if and only if A and B are isomorphic;

(II) The construction $\mathcal{Q}(\cdot)$ has a uniform cancellation. □

Note that (II) does not tell us that $\mathcal{Q}(\cdot)$ has an effective uniform cancellation procedure. If (I) is satisfied, all that our usual methods show is that there is an effective procedure which, given A and B, produces a Bij$(\overline{Q}A, \overline{Q}B)$-uniform cancellation. Unfortunately this falls short of an effective $\mathbf{FinSet}_{\overline{Q}}$-uniform cancellation. The point is that the Bij$(\overline{Q}A, \overline{Q}B)$-uniform cancellation we get may not extend to a $\mathbf{FinSet}_{\overline{Q}}$-uniform cancellation. For this reason, we are not in a position to replace "uniform cancellation" in (II) by "effective uniform cancellation procedure." It is not clear whether this difficulty is intrinsic or simply reflects the limitations of our approach.

9. Cancellation of Disjoint Union

In this and the following sections, we take A, B, and C to be non-empty sets.

The disjoint union $A \dot\cup B$ of two sets A and B can be modeled as $(\{1\} \times A) \cup (\{2\} \times B)$, although the details of this or any other representation are unimportant. The key points are that A and B need not be distinct, and that it is possible, given an element of the disjoint union of A and B, to recognize whether it came from A or from B. To emphasize the analogy with arithmetic addition, we will often write $A \dot\cup B$ as $A + B$.

Canceling the disjoint union construction $Q(X, C) = X + C$ is at once straightforward and important in practice.

Theorem 5 *For any fixed C, the construction $A \mapsto A + C$ can be effectively canceled.*

Proof Abstractly, the existence of a cancellation procedure follows from Theorem 2 and the simple fact that, for any finite group G and G-set C, two G-sets A and B are isomorphic if and only if G-sets $A + C$ and $B + C$ are isomorphic. Indeed, the effect of the functor $Q(\cdot, C)$ on the each of the G-set invariants σ, σ^+, τ (defined in Section 2) is to add a constant vector, clearly an invertible operation. (See [2].) Applying Theorem 2 we obtain the theorem. □

The following explicit cancellation procedure is well–known (though in the past it has been stated in the context of the complementary bijection principle; see below).

PROCEDURE: From a bijection $F : A + C \to B + C$, define a bijection $f : A \to B$ by setting $f(a)$ equal to the first element $b \in B$ obtained by iterating F on $a \in A$. •

Let us call this the *iterative procedure* for canceling disjoint union and write $\mathcal{I}F = f$. (The only *a priori* bound on the required number of iterations is $|C|$.) The iterative procedure is clearly effective since the construction of f uses only the information provided by the bijection F, not the properties or names of the elements of A, B or C.

Fred Richman has shown that even in this, the simplest instance of cancellation we can think of, the iterative procedure is not the unique effective cancellation procedure. For, suppose A, B, and C all have cardinality 2 or more. Define a cancellation procedure $\mathcal{J}$ (different from $\mathcal{I}$) as follows. Set $\mathcal{J}F = \mathcal{I}F$, unless F maps *just* two elements of A to C. If F maps just a_1 and a_2 to C, set $\mathcal{J}F(a_1) = \mathcal{I}F(a_2)$, $\mathcal{J}F(a_2) = \mathcal{I}F(a_1)$, and $\mathcal{J}F(a) = \mathcal{I}F(a)$ for a not equal to a_1 or a_2. In the same fashion, when $|A| = |B| \geq 2$ and $|C| \geq |A| - 2$, define a variant cancellation by modifying $\mathcal{I}$ on bijections F which map *all but* two elements of A to C. (At the end of this section, we settle the issue of what happens when one or more of A, B, C have small cardinality.)

It is evident that this cancellation principle is equivalent to the complementary bijection principle ([5], [11]); the only difference is that in our principle, the two copies of C in $A + C$ and $B + C$ are identified with one another, whereas in Gordon's they are not. Less trivially, we will now show that the involution principle of Garsia and Milne is also cancellation of disjoint union, in yet another guise.

The involution principle takes as given two "signed sets" $A = A_+ \dot\cup A_-$ and $B = B_+ \dot\cup B_-$, a sign-preserving bijection $f : A \to B$, and two involutions $i_A : A \to A$ and $i_B : B \to B$, such that all the fixed elements of i_A and i_B belong to the positive sets (A_+ and B_+) and i_A and i_B reverse the sign of elements they do not fix. The procedure of Garsia and Milne then finds a bijection between the fixed point sets A^{fix} and B^{fix} of the two involutions.

All we need do is switch the negative parts of the two signed sets. In more detail, form sets $D = A_+ \dot\cup B_-$ and $E = B_+ \dot\cup A_-$. Define a bijection $\tilde f : D \to E$ by $\tilde f|_{A_+} = f|_{A_+}$ and $\tilde f|_{B_-} = f^{-1}|_{B_-}$. Let A_+^{free} be the set of positive points *not* fixed by i_A, and define B_+^{free} similarly. Define a bijection $\tilde f_0 : A_+^{\text{free}} \dot\cup B_- \to B_+^{\text{free}} \dot\cup A_-$ by $\tilde f_0|_{A_+^{\text{free}}} = i_A|_{A_+^{\text{free}}}$ and $\tilde f_0|_{B_-} = i_B|_{B_-}$.

Now identify the sets $A_+^{\text{free}} \dot\cup B_-$ and $B_+^{\text{free}} \dot\cup A_-$ along the bijection $\tilde f_0$ and call the result C. Then $\tilde f$ is a bijection between $A^{\text{fix}} \dot\cup C$ and $B^{\text{fix}} \dot\cup C$.

The following theorem delineates when effective cancellation procedure for disjoint union is unique. This result will not be used again and may be

skipped on a first reading.

Theorem 6 *The iterative procedure is the unique effective cancellation procedure for disjoint union iff*

$|A| = |B| < 2$*; or*
$|C| = 1$ *and* $|A| = |B| \geq 4$*; or*
$|C| = 0$ *and* $|A| = |B| \neq 2$*.*

Proof Uniqueness is trivial when $|A| = |B| = 0$ or 1. Uniqueness is impossible when $|A| = |B| = 2$ because every effective procedure (for *any* cancellation problem) has a distinct "opposite" obtained by post-composing every output f with the nontrivial involution on B. When $|A| = |B| \geq 2$ and $|C| \geq 2$, we explained how to define a distinct cancellation procedure in Section 9.

To finish the proof we need two lemmas:

Lemma 6 *Suppose that an effective cancellation procedure takes* $F : A + C \to B + C$ *to* $f : A \to B$*. Suppose moreover that* $|A \cap F^{-1}(B)| \geq 2$*. Then* $f(A \cap F^{-1}(B)) = F(A) \cap B$*.*

Proof First note that $F(A \cap F^{-1}(B)) = F(A) \cap B$, so $|A \cap F^{-1}(B)| = |F(A) \cap B|$. Thus is suffices to prove $f(A \cap F^{-1}(B)) \subseteq F(A) \cap B$. By hypothesis, there are distinct elements $a_i \in A$, $i = 1, 2$ such that $F(a_i) = b_i \in B$. Suppose $f(a_1) \notin F(A) \cap B$. Writing elements of Sym A and Sym B in cycle notation, we have $((a_1 a_2), (b_1 b_2))F = F$, but $((a_1 a_2), (b_1 b_2))f \neq f$ (because $(((a_1 a_2), (b_1 b_2))f)(a_2) = f(a_1)$), violating equivariance. □

Lemma 7 *Suppose that an effective cancellation procedure takes* $F : A + C \to B + C$ *to* $f : A \to B$*. Suppose moreover that* $|A \cap F^{-1}(B)| \geq 3$*. Then* $f|_{A \cap F^{-1}(B)} = F|_{A \cap F^{-1}(B)}$*.*

Proof Lemma 6 shows that $(\overline{f}^{-1}F)(A \cap F^{-1}(B)) = A \cap F^{-1}(B)$. Let $j = \overline{f}^{-1}F)|_{A \cap F^{-1}(B)}$, and suppose j is not the identity map on $A \cap F^{-1}(B)$. Then there is some bijection $\rho : A \to A$ where $\rho(a) = a$ for $a \notin A \cap F^{-1}B$

such that $\rho|_{A\cap F^{-1}B}$ does not commute with j (because $|A \cap F^{-1}(B)| \geq 3$). Then $(\rho, F\rho F^{-1})F = F$, but $(\rho, F\rho F^{-1})f \neq f$ even on $A \cap F^{-1}B$, violating equivariance. □

Now all that remains is to prove is that the cancellation procedure is unique when $|A| = |B| \geq 3$ and $|C| = 0$ and when $|A| = |B| \geq 4$ and $|C| = 1$. When $|A| = |B| \geq 3$ and $|C| = 0$, Lemma 7 implies $f = F$. When $|A| = |B| \geq 4$ and $|C| = 1$, then f and F differ on at most a single element of A, so there is no choice about f in this case either. □

10. Noncancellation of Products

Now we consider the Cartesian product construction $X \mapsto X \times C$. Fix a set S, $|S| > 1$, and set both A and C equal to the set of linear orderings of S. Let B be the set of permutations of S. We may define a bijection between $A \times C$ and $B \times C$ by using the fact that two linear orderings of S induce a permutation of S. There can be no canonical bijection between A and B, however: the set of permutations of S has a distinguished element, namely the identity permutation, but the set of linear orderings does not.

More generally:

Theorem 7 *For any fixed C of cardinality greater than 1, the construction $A \mapsto A \times C$ cannot be effectively canceled.*

Proof Let G be a group of order $|C|$, which we identify with C. Let G_t be the G-set obtained by having G act on itself by translation (so that g sends g' to gg') and G_c be the G-set obtained by having G act on itself by conjugation (so that g sends g' to $gg'g^{-1}$). G_t and G_c cannot be isomorphic, as G_t is transitive and G_c is not. Nevertheless $h : G_t \times G_t \to G_c \times G_t$ defined by $h(g_1, g_2) = (g_1 g_2^{-1}, g_2)$ is always an isomorphism of G-sets. The example in the first paragraph is essentially the special case where G is a symmetric group. (See Lemma 1 of [1] for a statement of what is essentially the same idea, applied in a different context.) □

11. Cancellation of Pointed Products

The situation is altogether different when we multiply A and B by a set C with a distinguished point. As is customary, we call a set with a distinguished point a *pointed set*; the operation of forming the product of a (not necessarily pointed) set with a fixed pointed set will be called a *pointed product* construction.

Theorem 8 *For any fixed C with a distinguished point, the construction $A \mapsto A \times C$ can be effectively canceled.*

Proof From the group-theoretic viewpoint, the analogue of such a C is a G-set with a fixed point. Thus, we must show that if A and B are G-sets and C is a G-set with a fixed point, and $A \times C$ is isomorphic to $B \times C$, then A is isomorphic to B.

A G-set A is determined up to isomorphism by the invariant, σ_A^+, the integer-function on the set of subgroups of G which records how many elements are fixed by each subgroup H. The element (a, c) of $A \times C$ is fixed by H if and only if both a and c are fixed by H. Thus, for each H, $\sigma_A^+(H)\sigma_C^+(H) = \sigma_{A\times C}^+(H) = \sigma_{B\times C}^+(H) = \sigma_A^+(H)\sigma_C^+(H)$. Since C contains a fixed point, $\sigma_C^+(H) > 0$. It follows that $\sigma_A^+(H) = \sigma_A^+(H)$ for all H, so that A and B are isomorphic as G-sets. From Theorem 2 it now follows that there is a cancellation procedure for multiplication of sets by the pointed set C. □

We may also exhibit this cancellation principle more explicitly by the following construction. A and B are finite sets, and C is a finite pointed set with a distinguished element called $*$. Let $F : A \times C \to B \times C$ be a bijection.

PROCEDURE: We define a map $F_* : A \to B$ by letting $F_*(a)$ be the projection of $F((a, *))$ onto B. A map $F_*^{-1} (= (F^{-1})_*) : B \to A$ is defined similarly. Iteration of the map $F_* \cup F_*^{-1} : A \cup B \to A \cup B$ produces some cycles. We then use F_* to pair any element a of A that occurs in a cycle with the element $F(a)$. This gives us a nontrivial partial bijection between A and B, say from $\tilde{A}$ to $\tilde{B}$. Multiplying by C we get a nontrivial partial bijection from $A \times C$ to $B \times C$ taking $\tilde{A} \times C$ to $\tilde{B} \times C$. This induces a bijection from $(A \setminus \tilde{A}) \times C$ to $(B \setminus \tilde{B}) \times C$ (see Section 9) and now we may iterate the entire process until we get a bijection from A to B. •

It should be noted that this construction was only discovered after the group-theoretical approach suggested its feasibility. It is also worth verifying that for the example given in the first paragraph of Section 10, if we choose a particular linear ordering of the set S to serve as the distinguished point, then the construction we have just described induces the standard bijection between linear orderings of S and permutations of S, relative to that specified linear ordering.

Notice the contrast with the situation when C has no distinguished point. Let G be Aut C, and assume that the action of G on C has no fixed points. Think of the data $\sigma_A^+(H)$ and $\sigma_{A\times C}^+(H)$ (with H varying) as forming vectors, and think of the data $\sigma_C^+(H)$ as forming the entries of a diagonal matrix M, so that $M\sigma_A^+ = \sigma_{A\times C}^+$. Since the diagonal element $\sigma_C^+(G)$ of the diagonal matrix M is zero, M is singular, and we can find two G-sets A and B such that the vectors σ_A^+ and σ_B^+ differ by some non-zero vector in the null-space of M. Then $A \times C$ and $B \times C$ are isomorphic as G-sets, even though A and B are not. Thus, if C is some structured set, then Cartesian product by C can be canceled if and only if C has a distinguished point as part of its structure.

Problem 4 *We have just seen that if C is a structured finite set and A and B are arbitrary finite sets, then a bijection between $A \times C$ and $B \times C$ does not in general yield a bijection between A and B unless the structure on C singles out a particular point. However, our proof was algebraic, and not combinatorial. Can one give a more combinatorial understanding of this fact? More specifically, it would be desirable to have some generalization of the counterexamples from Section 10 that applied not just to G-sets $C = G_t$ (with $|C| > 1$) on which G acts freely, but to all G-sets that have no fixed point.*

A variant of the preceding procedure for canceling Cartesian product by a pointed set yields an explicit cancellation procedure for the example mentioned in connection with the Equivariance Criterion for Bifunctors, so we include it here. Suppose we have a bijection

$$F : (A \times A) + (B \times B) \to (A \times B) + (B \times A).$$

We can reconstruct from this a bijection between A and B as follows.

PROCEDURE: If A and B are empty, there is nothing to do. Otherwise, for $a \in A$, $F((a,a))$ is either a pair $(a',b) \in A \times B$ or a pair $(b,a') \in B \times A$. Either way, define a map $\alpha : A \to B$ by $\alpha(a) = b$. Define a map $\beta : B \to A$ similarly. As in the procedure that implemented Theorem 8, use these two maps to define a bijection between two non-empty subsets A_0 and B_0 of A and B, respectively. Write $A = A_0 + A_1$, $B = B_0 + B_1$. The bijection between A_0 and B_0 gives us a bijection between $(A_0 \times A_0) + (B_0 \times B_0)$ and $(A_0 \times B_0) + (B_0 \times A_0)$. It also gives a bijection between $(A_0 \times A_1) + (A_1 \times A_0) + (B_0 \times B_1) + (B_1 \times B_0)$ and $(B_0 \times A_1) + (A_1 \times B_0) + (A_0 \times B_1) + (B_1 \times A_0)$. Subtracting these bijections from the bijection between $(A \times A) + (B \times B)$ and $(A \times B) + (B \times A)$ (again, see Section 9), we obtain a bijection between $(A_1 \times A_1) + (B_1 \times B_1)$ and $(A_1 \times B_1) + (B_1 \times A_1)$. Repeat as needed. •

It is worth pointing out, as an aside, that the preceding procedure cannot be generalized so as to apply to the situation in which one is given a bijection between $A \times C + B \times D$ and $A \times D + B \times C$ and one wishes to obtain *either* a bijection between A and B *or* a bijection between C and D; for, if one lets G be the cyclic group of order 6 and lets A, B, C, D be G-sets which are composed of 1 orbit of size 6, 2 orbits of size 3, 1 orbit of size 6, and 3 orbit of size 2, respectively, then it is easy to check that A and B are non-isomorphic and C and D are non-isomorphic yet $A \times C + B \times D$ and $A \times D + B \times C$ are isomorphic.

Another sort of variant of cancellation of pointed products is what one might call "relative cancellation." Given a bijection between $A \times C \times D$ and $B \times C \times D$ and a map from C to D, one can construct a bijection between $A \times C$ and $B \times C$. (Note that in the case where C has a single point, so that the map from C to D is nothing more than the choice of a distinguished element of D, this construction cancels a bijection from $A \times D$ to $B \times D$ to obtain a bijection from A to B, as in Theorem 8.

PROCEDURE: Given an element of $A \times C$, we can map it into $A \times C \times D$ (using our map from C to D), apply the given bijection to get an element of $B \times C \times D$, and then project back to $B \times C$. Similarly, we can define a function from $B \times C$ to $A \times C$. This gives a partial bijection between $A \times C$ and $B \times C$. We can lift this to a partial bijection between $A \times C \times D$ and $B \times C \times D$ that acts as the identity on the third coordinate. Call the subsets of $A \times C$ and $B \times C$ that are in the partial bijection E and F. Then our lifted

bijection between $E \times D$ and $F \times D$ can be "subtracted" from our original bijection between $A \times C \times D$ and $B \times C \times D$ (in a fashion that by now should be familiar) to yield a bijection that matches elements of $(A \times C \times D) \setminus (E \times D)$ with elements of $(B \times C \times D) \setminus (F \times D)$. Now we are in a position to iterate the procedure. •

Note that this construction can be applied in the case where $D = C^2$, since we have the diagonal embedding of C into its Cartesian square.

One is tempted to link in one's mind the fact that $|A \times C| = |B \times C|$ does not necessarily imply $|A| = |B|$ (C could be empty) and the fact that a bijection between $A \times C$ and $B \times C$ does not necessarily yield a bijection between A and B (one needs to have a distinguished element of C). However, if the reader examines our argument carefully, it will be seen that no such link can be found there. In the first place, our method considers cancellation problems "one cardinality at a time"; in no case have we attempted to describe situations in which cardinalities of sets are unknown to either Programmer or Machine. In the second place, even when C is empty, our reason for Machine's inability to generate a bijection between A and B given a bijection between $A \times C$ and $B \times C$ is not that the cardinalities of A and B might be different (they are known to be the same, by both Programmer and Machine), but rather that the bijection between $A \times C$ and $B \times C$ (empty sets both) gives absolutely no information, and thus cannot break the pre-existing symmetry of the situation. Still, one can't help feeling that the need for a distinguished element is in some sense an "effective version" of the hypothesis that C is non-empty, and is inclined to wonder whether some meta-principle that would explain this coincidence is waiting to be discovered.

Cartesian product by a pointed set cannot generally be canceled anonymously. To see this, let A be a finite subset of the real line such that A is symmetrical about the origin and contains an odd number of positive elements. Let C be a finite subset of the real line such that C is symmetrical about the origin and does not contain 0. Define the bijection $F : A \times C \to A \times C$ by $F(a, c) = (ac/|c|, -c)$. If $\alpha : A \to A$ is multiplication by -1 and $\beta : A \to A$ is the identity map, then $\overline{\alpha} = F\overline{\beta}F$. Because the permutation α has odd parity, we cannot have $\alpha \neq f\beta f$ for any bijection $f : A \to A$. Hence, taking G to be the trivial group, we may appeal to Theorem 3.

We conclude this section with some open problems.

Problem 5 *Suppose A, B, C, D are finite sets with $|A| = |B|$ and $|C| = |D|$, and suppose we have a bijection between $A^2 + 4 \times C \times D + B^2$ and $2C^2 + 2 \times A \times B + 2 \times D^2$, where "2" and "4" respectively denote some fixed sets of 2 and 4 (mutually distinguished) elements. The General Equivariance Criterion of Section 14 may be used to show that we can cancel this bijection to obtain a bijection between A and B and a bijection between C and D. Is there a polynomial-time algorithm which will achieve this? Note that the given bijection embodies the algebraic relation $(|A|-|B|)^2 = 2(|C|-|D|)^2$, and that the desired bijections embody the relations $|A| - |B| = 0$ and $|C| - |D| = 0$; consequently, the sought-for cancellation principle could be construed as a combinatorial representation of the irrationality of the square root of 2.*

Problem 6 *If C is a pointed set, does an injection from $A \times C$ to $B \times C$ determine an injection from A to B?*

Problem 7 *If C is a pointed set, does a surjection from $A \times C$ to $B \times C$ determine a surjection from A to B?*

12. Cancellation of Powers

Theorem 9 *For any positive integer m, the construction $A \mapsto A^m$ (the mth Cartesian power of A) can be effectively canceled.*

Proof Let G be a finite group, and let A and B be G-sets such that the constructed G-sets A^m and B^m are isomorphic. For all subgroups H of G, $(\sigma_A^+(H))^m = \sigma_{A^m}^+(H) = \sigma_{B^m}^+(H) = (\sigma_B^+(H))^m$, implying $\sigma_A^+(H) = \sigma_B^+(H)$. It follows that A and B are isomorphic G-sets. Theorem 2 then permits us to deduce the theorem. □

A cancellation procedure that generates a bijection in time polynomial in $|A|$ and $|B|$ eluded us, but fortunately, Aaron Meyerowitz and Fred Richman saw an early version of this paper and came up with the following.

Let $F : A^m \to B^m$ be a bijection, and let $f_0 : A_0 \to B_0$ be a bijection between (possibly empty) sets $A_0 \subseteq A$, $B_0 \subseteq B$.

PROCEDURE: f_0 induces a bijection between A_0^m and B_0^m, so using our procedure for cancellation of disjoint union, we may derive from F and f_0 a bijection $F' : A^m \setminus A_0^m \to B^m \setminus B_0^m$. We define a map $F_* : A \setminus A_0 \to B \setminus B_0$ by letting $F_*(a)$ be the first component of the m-tuple $F'(a, a, ..., a)$ that is not in B_0. A map $F_*^{-1} : B \setminus B_0 \to A \setminus A_0$ is defined similarly. Proceeding as we did in Section 11, we may use F_* and F_*^{-1} to get a bijection between a non-empty subset of $A \setminus A_0$ and a non-empty subset of $B \setminus B_0$. This permits us to extend our partial bijection f_0 to a more inclusive partial bijection. Iterating this procedure, we will eventually arrive at a full bijection f between A and B. •

Problem 8 *It is easy to see that the method of proof of Theorem 9 shows more generally that a combinatorial construction that sends A to a disjoint union of Cartesian powers of A (where A^0 is just a single point) can be effectively canceled. Is there a general scheme for devising polynomial-time procedures that implement cancellation?*

One consequence of Theorem 9 is that a bijection between $A \times C$ and $B \times C$ *does* determine a bijection between A and B in the case $C = A + B$. For, a bijection between $A \times (A + B)$ and $B \times (A + B)$ yields a bijection between $(A \times A) + (A \times B)$ and $(B \times A) + (B \times B)$; if we identify $A \times B$ with $B \times A$ in the obvious way and apply cancellation of disjoint unions, we get a bijection between $A \times A$ and $B \times B$, and cancellation of powers gives us a bijection between A and B. In contrast, in the case where C is A (or B), a bijection between $A \times C$ and $B \times C$ does *not* necessarily give us a bijection between A and B, as the example given at the beginning of Section 10 shows.

Problem 9 *For what finite sets C defined in terms of A and B via $+$ and $\times$ is it the case that every bijection $A \times C \to B \times C$ can be canceled?*

A simple example shows why anonymous cancellation (as in Corollary 1 and Theorem 3) is generally impossible for the Cartesian square construction. Let $A = \{-1, 1\}$ and regard A^2 as a subset of the complex plane, that is, associate (j, k) with $j + k\sqrt{-1}$. Let $F : A^2 \to A^2$ be multiplication by $\sqrt{-1}$. Let $\alpha : A \to A$ be multiplication by -1 and $\beta : A \to A$ be the identity map. Then $\overline{\alpha} = F\overline{\beta}F$, but $\alpha \neq f\beta f$ for any bijection $f : A \to A$ (by a

parity argument, as in Section 11). Then taking G to be the trivial group, we appeal to Theorem 3. Moreover, A may be replaced by any finite set of complex numbers, symmetrical about the origin, in which the number of non-zero elements is congruent to 2 modulo 4.

The cancellation principles that have been illustrated in the article up to this point provide a useful tool-kit for the construction of many sorts of bijections, in a way that does not require one to get one's hands dirty in details. We give two examples that we find amusing.

First: Given a bijection between $A \times A \times C$ and $B \times B \times C$, one can construct a bijection between $A \times C$ and $B \times C$. For, multiplying the bijection by the identity map on C, one gets a bijection between $A \times A \times C \times C \cong (A \times C)^2$ and $B \times B \times C \times C \cong (B \times C)^2$. Now one can apply cancellation of Cartesian squares.

Second: Given bijections between $A^2 \to B \times C$, $B^2 \to A \times C$, and $C^2 \to A \times B$, one can construct bijections between A, B, and C. We have a bijection between $(A^2)^2 = A^2 \times A \times A$ and $(B \times C)^2 \cong B^2 \times C^2 \cong (A \times C) \times (A \times B) \cong B \times C \times A \times A$. Applying the relative cancellation procedure discussed in Section 11, one gets a bijection between $A^2 \times A$ and $B \times C \times A$. Putting this in a different way, we get a bijection between A^3 and $A \times B \times C$. But in exactly the same way we can get a bijection between B^3 and $A \times B \times C$. Thus we get a bijection between A^3 and B^3, which yields a bijection between A and B. A bijection between B and C can be obtained in the same way, and composing these bijections one gets a compatible bijection between A and C.

13. Noncancellation of Exponentials

The power set of S will be denoted 2^S. We show that a bijection between 2^{S_1} and 2^{S_2} does not generally induce a canonical bijection between S_1 and S_2 by finding a finite group G and a pair of G-sets S_1 and S_2 which are not isomorphic even though 2^{S_1} and 2^{S_2} are isomorphic. The counterexample is then the isomorphism between 2^{S_1} and 2^{S_2}, considered simply as a bijection between power sets. In this light, the G-action appears as relabeling-symmetries. If there were a canonical induced bijection between S_1 and S_2 it would also possess the same relabeling-symmetries. This is impossible, since G has a

different action on the two sets.

Let $\mathcal{O}$ be a G-orbit of H-type, with $s \in \mathcal{O}$ having stabilizer H. Pairing elements gs of $\mathcal{O}$ with cosets gH of H gives a G-set isomorphism between $\mathcal{O}$ and G/H. If K is a subgroup of G, the following are clearly all equivalent:
(i) gs is in the same K-orbit as $g's$
(ii) there exists $k \in K$ such that $gs = kg's$
(iii) there exists $k \in K$ such that $gH = kg'H$
(iv) there exist $k \in K$, $h \in H$ such that $g = kg'h$
(v) $KgH = Kg'H$.
Let $o(K,H)$ be the number of K-orbits in $\mathcal{O}$. The equivalence of (i) and (v) implies that $o(K,H)$ is also the number of double cosets of the form KgH. For any H' conjugate to H, $\mathcal{O}$ is also of H' type, so $o(K,H') = o(K,H)$. As the number of double cosets KgH equals the number of double cosets HgK (note $(HgK)^{-1} = Kg^{-1}H$), $o(K,H) = o(H,K)$. So for any K' conjugate to K, we also have $o(K',H) = o(K,H)$.

The key formula is

$$\sigma^{+}_{2^S}(K) = 2^{\sum_H o(H,K)\tau_S(H)},$$

where H ranges over representatives of the conjugacy classes of subgroups of G. The left side is the number of subsets of S which are stabilized by K. A subset of S is stabilized by K exactly if it is the union of K-orbits of S. The number of K-orbits of S is $\sum_{H \subseteq G} o(H,K)\tau_S(H)$ since each G-orbit of H-type decomposes into $o(H,K)$ K-orbits.

Example 1 Take $G = Z/2Z \times Z/2Z$, the Klein 4-group. Since G is Abelian the issue of conjugacy of subgroups is moot. The subgroups of G are G itself, the three 2-element subgroups R_1, R_2 and R_3, and $\{e\}$. Taking the subgroups in this order, the matrix

$$[o(H,K)] = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 1 & 1 & 2 \\ 1 & 1 & 2 & 1 & 2 \\ 1 & 1 & 1 & 2 & 2 \\ 1 & 2 & 2 & 2 & 4 \end{bmatrix}$$

is singular, with vector $[-2\ 1\ 1\ 1\ -1]$ in the kernel. The linear transformation $[o(H,K)]$ takes the same values on $[2\ 0\ 0\ 0\ 1]$ and $[0\ 1\ 1\ 1\ 0]$. Let S_1 be a

G-set with two orbits of type G and one of type $\{e\}$. Let S_2 be a G-set with one orbit of type R_i, $i = 1, 2, 3$. Then 2^{S_1} and 2^{S_2} are isomorphic G-sets.

Example 2 Take $G = S_3$, the symmetric group on three letters. The conjugacy classes of subgroups of G are represented G itself, the 3-element cyclic subgroup C_3, any of the 2-element subgroups R_i generated by a reflection, and $\{e\}$. Taking the subgroups in this order, the matrix

$$[o(H,K)] = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 1 & 2 \\ 1 & 1 & 2 & 3 \\ 1 & 2 & 3 & 6 \end{bmatrix}$$

is also singular, with vector $[-2\ 1\ 2\ -1]$ in the kernel, so the linear transformation $[o(H,K)]$ takes the same values at [2 0 0 1] and [0 1 2 0]. (Note that the subgroup-valued indices H, K in the matrix denoted by $[o(H,K)]$ range over a system of representatives of conjugacy classes of subgroups.) Let S_1 be a G-set with two orbits of type G and one of type $\{e\}$. Let S_2 be a G-set with one orbit of type C_3 and two of type R_i. Then 2^{S_1} and 2^{S_2} are isomorphic G-sets.

Thus we have:

Theorem 10 *The construction $A \mapsto 2^A$ cannot be effectively canceled.* □

The situation illustrated by the preceding examples is far from rare, as we will now show.

We thank Goetz Pfeiffer for a demonstration of the following

Theorem 11 *The matrix $[o(H,K)]$ has nonzero determinant precisely when G is cyclic.*

We have modified Pfeiffer's approach to avoid quoting results from representation theory. First we will need a

Lemma 8 *Let G be a finite group with subgroups H and K. Then the number of double cosets KgH is the scalar product of the permutation characters of H and K.*

Proof The number of double cosets KgH is the number of K-orbits in G/H. By the orbit counting formula, this is

$$\begin{aligned}\frac{1}{|K|}\sum_{k\in K}|\mathrm{Fix}_{G/H}k| &= \frac{1}{|K|}\sum_{k\in K}\frac{1}{|H|}\sum_{g\in G}\delta(kgH=gH)\\ &= \frac{1}{|K|}\frac{1}{|H|}\sum_{k\in K}\sum_{g\in G}\delta(g^{-1}kg\in H)\end{aligned}$$

where $\delta(\mathcal{P})$ is 1 if $\mathcal{P}$ is true and 0 if not. Since each pair $(k,g)\in K\times G$ satisfying $g^{-1}kg\in H$ gives rise to $|G|$ triples $(x,y,z)=(g'kg'^{-1},g',g'g)$ satisfying $y^{-1}xy\in K, z^{-1}xz\in H$, we may rewrite what we had as

$$\begin{aligned}&\frac{1}{|G|}\frac{1}{|K|}\frac{1}{|H|}\sum_{x\in G}\sum_{y\in G}\sum_{z\in G}\delta(y^{-1}xy\in K, z^{-1}xz\in H)\\ =\ &\frac{1}{|G|}\sum_{x\in G}\frac{1}{|K|}\sum_{y\in G}\delta(y^{-1}xy\in K)\frac{1}{|H|}\sum_{z\in G}\delta(z^{-1}xz\in H)\\ =\ &\frac{1}{|G|}\sum_{x\in G}\frac{1}{|K|}\sum_{y\in G}\delta(xyK=yK)\frac{1}{|H|}\sum_{z\in G}\delta(xzH=zH)\\ =\ &\frac{1}{|G|}\sum_{x\in G}|\mathrm{Fix}_{G/K}\,x|\ \ |\mathrm{Fix}_{G/H}\,x|\end{aligned}$$

as desired. □

W. Burnside's classic *Theory of Groups of Finite Order* introduces a matrix called there the "table of marks" associated to any finite group G. Let

$$G_1=\{e\},G_2,\ldots,G_s=G$$

be a sequence of representatives for conjugacy classes of subgroups of G ordered so that

$$|G_1|\le|G_2|\le\cdots\le|G_s|\ ,$$

and let O_i be a G-orbit of G_i-type, e.g., G/G_i. The "table of marks" is then the $s\times s$ matrix $B=[m_{ij}]$ where m_{ij} is just the number of elements of O_i fixed by G_j. The stabilizers of elements of O_i are conjugates of G_i, so $m_{ij}=0$ unless G_j is contained in some conjugate of G_i; in particular $m_{ij}=0$ if $i<j$. On the other hand m_{ii} is always positive since O_i certainly contains at least one G_i-stable element. Thus $[m_{ij}]$ is a lower triangular matrix with non-zero

entries on the diagonal, and the determinant of $[m_{ij}]$ does not vanish. Hence the columns of B are linearly independent.

Now let $C = [c_{in}]$ be the $s \times |G|$ matrix with rows indexed by the conjugacy classes of subgroups of G, columns indexed by the elements of G, and c_{in} equal to the number of elements fixed when g_n acts on O_i. The rows of C are by definition the permutations characters of G. The n^{th} column of C coincides with the column of B corresponding to the (conjugacy class of) the subgroup $\langle g_n \rangle$ of G generated by g_n, since $\langle g_n \rangle$ and g_n leave the same elements fixed. The rank of C is thus the number of conjugacy classes of cyclic subgroups of G.

By Lemma 8, CC^T is the $s \times s$ matrix $D = [d_{ij}]$, with rows and columns indexed by conjugacy classes of subgroups of G and $d_{ij} = |G| o(G_i, G_j)$. The rank of D cannot be more than the rank of C, so for D to have full rank, every subgroup of G, including G itself, must necessarily be cyclic.

That G be cyclic is also sufficient for D to have full rank. Let $\tilde{C} = [\tilde{c}_{ij}]$ be the $s \times s$ matrix with rows indexed by the subgroups of G, columns indexed by a generator for each subgroup and $\tilde{c}_{ij}$ equal to the number of elements fixed when g_j acts on O_i times the number of generators of $\langle g_j \rangle$. Clearly $\tilde{C}\tilde{C}^T = CC^T = D$, since $\tilde{C}$ amounts to "collecting" the identical rows of C. On the other hand $\tilde{C}$ is invertible since its columns are nonzero multiples of the columns in the table of marks. □

Comment The previous argument gives a formula for the determinant of $D = D(n)$ when $G = \mathbf{Z}/\langle n \rangle$:

$$\mathrm{Det}(D) = \prod_{d|n} d\phi(d) \ .$$

Alternatively, the determinant of a tensor product is given by the formula

$$\mathrm{Det}(A \otimes B) = \mathrm{Det}(A)^{\mathrm{dim(B)}} \cdot \mathrm{Det}(B)^{\mathrm{dim(A)}} \ .$$

If $n = pq$ with p and q relatively prime, then essentially

$$D(n) = D(p) \otimes D(q) \ .$$

If p is prime, then

$$\mathrm{Det}(D(p^n)) = p^{p^{n-1}(p-1)} \cdot \mathrm{Det}(D(p^{n-1}))$$

and $\mathrm{Det}(D(p)) = p - 1$.

Theorem 11 guarantees a rich supply of counterexamples to the cancellation of the power set construction, but in a small way it has a positive aspect as well—it shows that a bijection between power sets that has only cyclic symmetry can in fact be canceled.

Note that the key formula above generalizes to

$$\sigma^+_{b^S}(K) = b^{\sum_H o(H,K)\tau_S(H)}, \text{ for } b \geq 2$$

(where we think of b as both a number and as the b-element set $\{0, 1, 2, ..., b-1\}$). So, if S_1 and S_2 are G-sets such that 2^{S_1} and 2^{S_2} are isomorphic, then b^{S_1} and b^{S_2} isomorphic as well, for each $b \geq 2$. Thus even an infinite family of bijections $F_b : b^A \to b^B$, $b \geq 2$, may still fail to effectively determine a bijection between A and B.

What is more, any bijection $F_b : b^A \to b^B$ ($b \geq 2$) itself effectively determines bijections $F_{b'} : b'^A \to b'^B$ for all $b' \geq 2$. Indeed, let $G = \mathrm{Stab}\ F_b$. Then even though, as G-sets, A and B may not be isomorphic, the G-sets b'^A and b'^B are isomorphic, and any G-set isomorphism $F_{b'} : b'^A \to b'^B$ is invariant under relabelings that preserve F_b.

Problem 10 *Is it possible to effectively construct an $F_{b'}$ from F_b in time polynomial in* $\max(|b^A|, |b'^A|)$*?*

We also ask:

Problem 11 *Does a bijection between 2^{2^A} and 2^{2^B} induce one between 2^A and 2^B?*

As a final example, consider the "self-map construction" $QA = A^A$, the set of maps from A to itself. We claim that this construction cannot be canceled. To describe a counterexample, let us say that a $\mathbf{Z}/\langle 2\rangle$-set with α fixed points and β 2-cycles ($\alpha + 2\beta$ elements in total) is of type α, β. Let A, B be of type 2,1 and 0,2, respectively. Then A^A (the set of functions from A to A) is isomorphic to B^B, since both are of type 32,8. This example is typical, in that for many simple constructions Q for which cancellation principles

do not exist, one may find counterexamples using just the 2-element group $G = \mathbf{Z}/\langle 2 \rangle$.

14. New Structures from Old

The variations given in Section 5 seem to point the way toward increasingly baroque forms of the Basic Equivariance Criterion. However, there is in fact a unified point of view which subsumes all of them, and which in some ways is simpler than any of them. The goal of this section is to explain this viewpoint, and to provide the setting in which all these variations can be conveniently proved.

The key insight is that a bijection between sets A and B is itself a kind of combinatorial structure on the set $A \dot\cup B$: a matching on the set $A \dot\cup B$ that pairs elements of A with elements of B. The situation we are in before we are given a specific bijection between A and B is also a structure on $A \dot\cup B$, specifically, a splitting of $A \dot\cup B$ into two parts of equal cardinality (A and B). For simplicity, we will assume for now that the two components of the partition are distinguishable from one another, though later in this section, when we develop these ideas formally, we will want to drop this assumption, since we want our framework to be able to deal with anonymous cancellation.

From this new point of view, questions about cancellation can be phrased so that they do not refer to bijections at all, but are merely questions about the relative specifiability of certain kinds of structures. For purposes of illustration, let us take a look at the the cancellation problem for Cartesian squares of sets from the new point of view. We are given, first of all, a splitting of the set $A \dot\cup B$, along with a splitting of the set $A^2 \dot\cup B^2$; these two splittings are not unrelated, but must be consonant with one another, in a way that will be formalized as a kind of functoriality. We are also given a matching of the set $A^2 \dot\cup B^2$ which is consonant with the splitting of $A^2 \dot\cup B^2$ (again, in a sense that will amount to nothing more than functoriality, once the right definitions are in place). What we seek is a matching of $A \dot\cup B$ which is consonant with our splitting of $A \dot\cup B$.

Four kinds of structure play a role, namely, splittings and matchings of $A \dot\cup B$

and $A^2 \dot{\cup} B^2$. It turns out that all our results fit into a "four structures" framework. This, in turn, can be fit into a "two structures" framework, where all the given data (initially embodied in three structures) can be rolled up into a single composite structure. However, before we can demonstrate the approach, we must explain, at last, what we mean by a "structure".

Define a *finite structure type* to be a pair $(\mathbf{S}, F)$ consisting of a category $\mathbf{S}$ and a functor F from $\mathbf{S}$ to **Bij** satisfying
(i) If $F(s) = u$ and $\tau : u \to u'$ is a morphism in **Bij**, then there exists a unique object $\tau s \in \mathbf{S}$ and a unique $\mathbf{S}$-morphism $\sigma : s \to \tau s$ such that $F(\sigma)=\tau$;
(ii) For $u \in \mathbf{Bij}$, $F^{-1}(u)$ is a finite set.

One should think of F as a forgetful functor taking structured objects to their underlying sets. We say that $F^{-1}(u)$ is the set of $(\mathbf{S}, F)$-structures on u. Isomorphism of $(\mathbf{S}, F)$-structures is simply isomorphism in the category $\mathbf{S}$. By (i) above, there is a (Sym u)-action on $F^{-1}(u)$. (Indeed, the set $F^{-1}(u)$ is actually functorial in u.) We define Aut s to be the subgroup of Sym $F(s)$ that fixes s.

A finite structure type is *effectively presented* if in addition there is a coding for any $(\mathbf{S}, F)$-structure s on any u by a (not necessarily unique) string of elements of u, possibly along with some other supplemental symbols from some finite alphabet, such that:
(iii) all the strings that represent $(\mathbf{S}, F)$-structures on u can be effectively enumerated;
(iv) it can be effectively determined whether two strings represent the same $(\mathbf{S}, F)$-structure;
(v) if $\tau : u \to u'$ is a bijection, S is a string representing an $(\mathbf{S}, F)$-structure s on u and the string S' is obtained from S by substituting elements of u' for elements of u according to τ, then S' represents τs.

All finite structure types to be discussed henceforth are effectively presented.

If $F_{2,1} : \mathbf{S}_2 \to \mathbf{S}_1$ is a functor such that $F_2 = F_1 \circ F_{2,1}$, we say that the finite structure type $(\mathbf{S}_2, F_2)$ *refines* the finite structure type $(\mathbf{S}_1, F_1)$ *via* $F_{2,1}$. Note that Aut s_1 acts on $F_{2,1}^{-1}(s_1)$.

Definition: An *effective procedure* for defining a $\mathbf{S}_2$-structure in terms of an $\mathbf{S}_1$-structure s_1 is a fixed sequence of instructions for Programmer to send to

Machine, such that if these instructions are applied by Machine to any string that codes s_1, the result will be a string that codes an $\mathbf{S}_2$-structure s_2, such that s_2 depends neither upon which string that codes s_1 is used by Machine nor upon any arbitrary choices that may be made by Machine in the course of carrying out Programmer's instructions.

Let $(\mathbf{S}_1, F_1)$ and $(\mathbf{S}_2, F_2)$ be effectively presented finite structure types such that $(\mathbf{S}_2, F_2)$ refines $(\mathbf{S}_1, F_1)$ via $F_{2,1}$.

Theorem 12 (Fixed Point Criterion) *Let u be a finite set. Given an $(\mathbf{S}_1, F_1)$-structure s_1 on u, one may effectively define an $(\mathbf{S}_2, F_2)$-structure s_2 on u in terms of s_1 such that $F_{2,1}(s_2) = s_1$ if and only if the* (Aut s_1)*-action on $F_{2,1}^{-1}(s_1)$ has a fixed point.*

Proof Suppose s_2 satisfying $F_{2,1}(s_2) = s_1$ is effectively defined in terms of s_1, but with no reference to the names of the elements of u. Then for any $\sigma : u \to u$ which preserves s_1, we have $\sigma s_2 = s_2$. Thus s_2 is a fixed point of the (Aut s_1)-action on $F_{2,1}^{-1}(s_1)$.

Conversely, let us assume that the (Aut s_1)-action on $F_{2,1}^{-1}(s_1)$ has a fixed point. Set $n = |u|$ and $\mathbf{n} = \{1, 2, \ldots, n\}$.

Programmer begins by asking Machine to find the first string in the lexicographic order which represents an $(\mathbf{S}_1, F_1)$-structure s_1' on $\mathbf{n}$ isomorphic to the structure s_1 on u. Next programmer asks Machine to choose a a bijection $\alpha : \mathbf{n} \to u$ such that $\alpha s_1' = s_1$.

Note that α gives a well-ordering of u. The resulting lexicographic ordering on the set of strings of elements of u first allows Machine to pick a unique representing string for each element of $F_{2,1}^{-1}(s_1)$, and then induces an order on the chosen strings. Machine may thereby determine a well-ordering of $F_{2,1}^{-1}(s_1)$.

Finally, Programmer asks Machine to output as s_2 the lexicographically least (Aut s_1)-fixed point in $F_{2,1}^{-1}(s_1)$.

We must check that s_2 is independent of the choice of α. Suppose Machine had chosen instead $\alpha' : \mathbf{n} \to u$ such that $\alpha' s_1' = s_1$. Then $\alpha'\alpha^{-1} : u \to u$ belongs to Aut s_1. Now the well-ordering of $F_{2,1}^{-1}(s_1)$ determined by α' may

be obtained by pushing forward the well-ordering of $F_{2,1}^{-1}(s_1)$ determined by α along the action of $\alpha'\alpha^{-1}$ on $F_{2,1}^{-1}(s_1)$. Since $\alpha'\alpha^{-1} \in \mathrm{Aut}\ s_1$, the (Aut s_1)-fixed points in $F_{2,1}^{-1}(s_1)$ are assigned the same ordinal by the α and α' orderings. In particular, the least fixed points coincide, so s_2 is well-defined. □

Let $(\mathbf{S}, F_S)$, $(\tilde{\mathbf{S}}, F_{\tilde{S}})$, $(\mathbf{T}, F_T)$ and $(\tilde{\mathbf{T}}, F_{\tilde{T}})$, be effectively presented finite structure types. Suppose that $(\tilde{\mathbf{S}}, F_{\tilde{S}})$ refines $(\mathbf{S}, F_S)$ via a functor $F_{\tilde{S},S}$ and $(\tilde{\mathbf{T}}, F_{\tilde{T}})$ refines $(\mathbf{T}, F_T)$ via a functor $F_{\tilde{T},T}$. Let $Q : \mathbf{S} \to \mathbf{T}$ be a functor.

Theorem 13 (General Equivariance Criterion) *Having fixed an* $(\mathbf{S}, \mathbf{F_S})$*-structure* s *on set* u *one can effectively define an* $(\tilde{\mathbf{S}}, F_{\tilde{S}})$*-structure* $\tilde{s}$ *on* u *such that* $F_{\tilde{S},S}(\tilde{s}) = s$ *from any* $(\tilde{\mathbf{T}}, F_{\tilde{T}})$*-structure* $\tilde{t}$ *on* u *such that* $F_{\tilde{T},T}(\tilde{t}) = Qs$*, if and only if there is an* (Aut s)*-equivariant map* $F_{\tilde{T},T}^{-1}(Qs) \to F_{\tilde{S},S}^{-1}(s)$.

Proof We are going to derive this result from the Fixed Point Criterion. To apply the Fixed Point Criterion we must specify a finite set u and finite structure types $(\mathbf{S}_1, F_1)$ and $(\mathbf{S}_2, F_2)$, where $(\mathbf{S}_2, F_2)$ refines $(\mathbf{S}_1, F_1)$ via a functor $F_{2,1}$.

The u of the Fixed Point Criterion will coincide with our present u. The finite structure type $(\mathbf{S}_1, F_1)$ of the Fixed Point Criterion will be our present $(\mathbf{S}, F)$. Our goal now is to devise a suitable finite structure type $(\mathbf{S}_2, F_2)$.

To say that one can effectively define an $(\tilde{\mathbf{S}}, F_{\tilde{S}})$-structure $\tilde{s}$ on u such that $F_{\tilde{S},S}(\tilde{s}) = s$ from any $(\tilde{\mathbf{T}}, F_{\tilde{T}})$-structure $\tilde{t}$ on u such that $F_{\tilde{T},T}(\tilde{t}) = Qs$, is just to say that one can effectively define a map $m : F_{\tilde{T},T}^{-1}(Qs) \to F_{\tilde{S},S}^{-1}(s)$. (Maps m of this sort generalize the maps $\mathcal{F}_{A,B} : \mathrm{Bij}(QA, QB) \to \mathrm{Bij}(A, B)$ we studied earlier.) The idea is to regard such a map m as a structure on u refining s.

Now we are going to define the new finite structure type $(\mathbf{S}_2, F_2)$. The category $\mathbf{S}_2$ will have objects that are triples (m, s, u) where u is a finite set, s is an $(\mathbf{S}, F_S)$-structure on u and m is a map $m : F_{\tilde{T},T}^{-1}(Qs) \to F_{\tilde{S},S}^{-1}(s)$. A morphism $\nu : (m, s, u) \to (m', s', t')$ in $\mathbf{S}_2$ is given by a bijection $\tau : t \to t'$ such that $s' = \tau s$ and $m' = \tau m$. (The action of τ on m is induced by the actions of τ on $F_{\tilde{T},T}^{-1}(Qs)$ and $F_{\tilde{S},S}^{-1}(s)$.) Define the functor $F_2 : S_2 \to \mathbf{Bij}$ so that on objects $F_2(m, s, u) = u$ and on morphisms $F_2(\nu) = \tau$. Define the functor

$F_{2,1}$ so that on objects $F_{2,1}(m,s,u) = s$ and on morphisms $F_{2,1}(\nu) = \sigma$ where $\sigma : s \to s' = \tau s$ is the unique morphism such that $F_1(\sigma) = \tau$.

By construction, the (Aut s)-action on $F_{2,1}^{-1}(s)$ has a fixed point precisely when there exists an (Aut s)-equivariant map from $F_{\tilde{T},T}^{-1}(Qs)$ to $F_{\tilde{S},S}^{-1}(s)$. The result now follows from the Fixed Point Criterion. □

Two finite structure types, introduced informally earlier, will play a distinguished role in the study of bijections. The finite structure type **Split** will be the pair consisting of

> (i) the category whose objects are finite sets each carrying an equivalence relation that determines exactly two equivalence classes, these of equal size, and whose morphisms are bijections which preserve the equivalence relations;
> (ii) the functor from this category to **Bij** which forgets the equivalence relation.

The finite structure type **Match** is the refinement of **Split** where the objects additionally carry a bijection between the two equivalence classes which the morphisms preserve.

To obtain results about canceling bijections from the General Equivariance Criterion, just take $(\mathbf{S}, F_S)$ and $(\mathbf{T}, F_T)$ to be refinements of **Split** and $(\tilde{\mathbf{S}}, F_{\tilde{S}})$ and $(\tilde{\mathbf{T}}, F_{\tilde{T}})$ to be the corresponding refinements of **Match**.

Even with our interests centered on bijections, we may be led to consider correspondences more general than bijection, and so benefit from the flexibility of the General Equivariance Criterion.

Example By a *signed set* we mean a set A partitioned into a set A_+ of *positive* elements and a set A_- of *negative* elements. To reinforce this intuition, we will write such a signed set as $A_+ \dot{-} A_-$, rather than adopting some more accurate but unintuitive notation like (A_+, A_-). A should be thought of as a combinatorial instantiation of the integer $|A_+| - |A_-|$, which we denote by $||A||$. A *signjection* between signed sets $A = A_+ \dot{-} A_-$ and $B = B_+ \dot{-} B_-$ is just a bijection between $A_+ + B_-$ and $A_- + B_+$. (One may also regard a signjection as consisting of a sign-reversing partial matching of the elements of A, a sign-reversing partial matching of the elements of B, and a sign-preserving

bijection between the remaining unmatched elements of A and of B.) The set of signjections from A to B we write as $\mathrm{Sign}(A, B)$.

The natural way to multiply a set X by a signed set $Y = Y_+ \dot{-} Y_-$ (corresponding to multiplying a natural number by an integer, under our interpretation) is given by the functor Q with $Q(X) = (X \times Y_+) \dot{-} (X \times Y_-)$. Put $Y_+ = A$ and $Y_- = B$, with A, B fixed. As degenerate cases, we may take $X = A$ and $X = B$, respectively. It is easily proved that $||QA|| = ||QB||$ implies $|A| = |B|$; we seek a cancellation principle corresponding to this. (Since a signjection between QA and QB amounts to a bijection between $(A \times A) + (B \times B)$ and $(A \times B) + (B \times A)$, what we have is just a new viewpoint on the problem we posed in connection with the Equivariance Criterion for Bifunctors and solved in Section 11.) In this case the General Equivariance Criterion tells us that effective cancellation is possible if and only if there is a $(\mathrm{Sym}\, A \times \mathrm{Sym}\, B)$-equivariant map from $\mathrm{Sign}(QA, QB)$ to $\mathrm{Bij}(A, B)$. Specifically, the finite structure type $(\tilde{\mathbf{T}}, F_{\tilde{T}})$ in the criterion would consist of a pair of signed sets (playing the roles of QA and QB) and a signjection between them.

Problem 12 *More generally, define* $(X_+ \dot{-} X_-) \times (Y_+ \dot{-} Y_-)$ *as* $(X_+ \times Y_+ + X_- \times Y_-) \dot{-} (X_+ \times Y_- + X_- \times Y_+)$. *Now suppose* X, Y, *and* Z *are signed sets. Given a signjection between* $X^n + Y^n$ *and* Z^n *(with* $n \geq 3$*), can one effectively define a signjection between* $X \times Y \times Z$ *and the empty set?*

Example In Section 5 we discussed two ways of measuring the failure when cancellation is impossible. Alternatively, when effective cancellation for $\mathcal{Q}(\cdot, \cdot)$ applied to A, B and C fails, we could consider settling for less. For example, we could seek a non-empty partial bijection between A and B, perhaps of a certain size; we could seek some other sort of map from A to B; we could seek a (small, non-empty) set of bijections; we could seek a bijection between RA and RB, where R is some other construction. Suitably specialized, the General Equivariance Criterion may be brought to bear on any of these situations. All one need do is specify suitable finite structure types $(\tilde{\mathbf{S}}, F_{\tilde{S}})$ these alternative outputs.

We note in passing that small, non-empty, effectively definable *sets of* bijections are closely related to small $(\mathrm{Aut}\ s)$-orbits in $F_{2,1}^{-1}(s)$ as in the proof

of the General Equivariance Criterion, with the General Equivariance Criterion being the special case of orbits of size (i.e., fixed points). This seeming generalization of the General Equivariance Criterion does not actually give any added generality, since unions of orbits in a G-set (of which equivariant sets of bijections are just a special case) are tantamount to fixed points in a derived G-set, namely, the power set with the inherited G-action.

15. Categorical Viewpoint

Till now, we have regarded constructions Q as endofunctors of the category **Bij** of finite sets and bijections. It is better here to consider the functor $\widehat{Q}$ from **Bij** to $\mathbf{Bij}_Q$, the full image of Q, which is defined as follows. The objects of $\mathbf{Bij}_Q$ are the objects of **Bij**, but

$$\mathrm{Hom}_{\mathbf{Bij}_Q}(A,B) = \mathrm{Bij}(QA,QB) \ .$$

On objects $\widehat{Q}$ is the identity map (this avoids the technical annoyance that $QA = QB$ does not generally imply $A = B$). On maps $\widehat{Q}$ coincides with Q. Note that $\widehat{Q}$ is surjective on objects where Q is not.

Our notion of a canonical cancellation for a construction Q lies between two extremes. A certain non-constructive view might regard Q as cancelable simply if the cardinality of QA determines the cardinality of A. The purely functorial view says Q is cancelable if there is a functor from $\mathbf{Bij}_Q$ back to **Bij** which is a one-sided inverse to Q. In general one cannot expect such functors to exist. Indeed when $|QA| > |A| > 4$, the group homomorphism

$$\mathrm{Hom}_{\mathbf{Bij}_Q}(A,A) = \mathrm{Bij}(QA,QA) \to \mathrm{Bij}(A,A)$$

has image of size at most 2 since the symmetric group $\mathrm{Hom}_{\mathbf{Bij}_Q}(A,A)$ contains a simple subgroup of index 2. In particular, the homomorphism cannot be surjective.

The question arises then, if our equivariant maps

$$\mathcal{F}_{A,B} : \mathrm{Hom}_{\mathbf{Bij}_Q}(A,B) \to \mathrm{Bij}(A,B)$$

do not define functors, how close do they come? Indeed we have functor-like gadgets $\mathcal{F}$, mapping objects to objects and arrows to arrows from $\mathbf{Bij}_Q$ to

Bij. Let us say that $\mathcal{F}$ is a *semifunctorial cancellation* for $\widehat{Q}$ if for any pair of composable arrows x and y,

$$\mathcal{F}(xy) = \mathcal{F}(x)\mathcal{F}(y)$$

provided at least one of x or y is in the image of the functor $\widehat{Q}$. Effective cancellation procedures give semifunctorial cancellations, because they do not depend on the names of elements, and morphisms in the image of $\widehat{Q}$ are essentially those induced by relabeling.

Note that a single equivariant bijection $\mathcal{F}_{A,B}$ induces a semifunctorial cancellation on a component of the category; there are no further compatibility constraints. If we take $A = B$, we are just looking at maps from $\mathrm{Hom}_{\mathbf{Bij}_Q}(A, A) = \mathrm{Bij}(QA, QA)$ to $\mathrm{Bij}(A, A)$ as two-sided $\mathrm{Bij}(A, A)$-sets, so we are back to group theory. Since these are both symmetric groups, we frame the following:

Problem 13 *Classify pairs (G, H), with G and H symmetric groups satisfying $H \subseteq G$, according to the existence of a map $G \to H$ that respects G and H as two-sided H sets.*

Furthermore, we propose:

Problem 14 *Explore the notion of semifunctorial cancellation developed in Section 15 for functors outside of combinatorics.*

16. Foundational Implications

In a series of papers from the 1920's, Alfred Tarski and Adolf Lindenbaum worked on the arithmetic of infinite cardinals from a point of view somewhat similar to ours; see [8].

Some of our theorems may be cast as independence results about **ZF**, set theory without the Axiom of Choice. Let $(A_i)_{i \in \mathbf{N}}$ and $(B_i)_{i \in \mathbf{N}}$ be two sequences of finite sets indexed by the natural numbers, such that that $|A_i| = |B_i|$ for every i. We say that the sequences (A_i) and (B_i) are *isomorphic* if there

is a sequence $(f_i)_{i\in\mathbf{N}}$ of bijections $f_i : A_i \to B_i$. Such isomorphism is a non-trivial issue absent the Axiom of Choice. Our results show that if the sequences (A_i^2) and (B_i^2) are isomorphic, then so must be the sequences (A_i) and (B_i). On the other hand, our results and some routine forcing yield models where (2^{A_i}) and (2^{B_i}) are isomorphic, but (A_i) and (B_i) are not.

To each finite structure type $(\mathbf{S}, F)$ there is associated an **Axiom of Choice for $(\mathbf{S}, F)$-structures**: any family of non-empty sets, each of which carries a fixed $(\mathbf{S}, F)$-structure, has non-empty product. It is natural to study the logical implications between these axioms, and these may be quite subtle even for the most basic sorts of structures, as in [3].

Our positive results have direct implications for such a study, but not our negative results. For example, let an $(\mathbf{S_1}, F_1)$-structure on a finite set C consist of sets A and B and a bijection between C and $\mathrm{Bij}(A^2, B^2)$. Let an $(\mathbf{S_2}, F_2)$-structure on C consist of sets A and B and a bijection between C and $\mathrm{Bij}(A, B)$. Theorem 9 shows directly that the **Axiom of Choice for $(\mathbf{S_1}, F_1)$-structures** implies the **Axiom of Choice for $(\mathbf{S_2}, F_2)$-structures**. On the other hand, let an $(\mathbf{S_3}, F_3)$-structure on a finite set C consist of sets A and B and a bijection between C and $\mathrm{Bij}(2^A, 2^B)$. Then we have the

Problem 15 *Does the Axiom of Choice for $(\mathbf{S_3}, F_3)$-structures imply the Axiom of Choice for $(\mathbf{S_2}, F_2)$-structures?*

A negative answer here would be stronger than the main result of Section 13, which says only that a choice function for the *particular* family $\mathrm{Bij}(2^{A_i}, 2^{B_i})$ does not generally suffice to define a choice function for the family $\mathrm{Bij}(A_i, B_i)$.

Ineffective arguments may play a role in this new context. To illustrate how this could happen, we take leave of our combinatorial context and we consider an example concerning an Axiom of Choice for a structure type on sets that are not necessarily finite. We will show that the **Axiom of Choice for compact Hausdorff spaces** does not imply the full **Axiom of Choice.**

The axiom **Weak Tychonoff** says any product of compact *Hausdorff* spaces is compact (where compact means *any open cover has a finite subcover*). It is known that Weak Tychonoff does not imply the full **Axiom of Choice**.

We show that Weak Tychonoff implies the **Axiom of Choice for compact**

Hausdorff spaces. If X_i, $i \in I$ is a family of compact Hausdorff spaces, and P is a one-point space, Weak Tychonoff assures us that $X = \prod_{i\in I}(X_i \dot\cup P)$ is compact. For any finite subset J of I, set $X_J = (\prod_{i\in I\setminus J}(X_i \dot\cup P)) \times (\prod_{i\in J} X_i)$. The X_J form a family of closed sets with the finite intersection property. Thus $\bigcap_J X_J = \prod_{i\in I} X_i$ is non-empty.

Even had we taken Weak Tychonoff in the effective form that fixes, once and for all, a choice of finite subcover for *every* open cover of *every* product of compact Hausdorff spaces, this argument still fails to define a point in $\prod_{i\in I} X_i$.

We thank Fred Linton for useful discussions, and Aaron Meyerowitz for his extremely careful reading of the manuscript.

REFERENCES

1. R. Appleson and L. Lovász, A characterization of cancelable k-ary structures, *Period. Math. Hung.* **6** (1975), 17-19.

2. W. Burnside, "Theory of Groups of Finite Order," Dover, New York, 1955.

3. J. H. Conway, Effective implications between the "finite" choice axioms, *in* "Cambridge Summer School in Mathematical Logic," Lecture Notes in Mathematics no. 337, Springer-Verlag, New York, 1973.

4. A. M. Garsia and S. C. Milne, A Rogers-Ramanujan bijection, *J. Comb. Th. Ser. A* **31** (1981), 289-339.

5. B. Gordon, Sieve-equivalence and explicit bijections, *J. Comb. Th. Ser. A* **34** (1983), 90-93.

6. N. Jacobson, "Basic Algebra I," W. H. Freeman, San Francisco, 1974.

7. A. Joyal, Une théorie combinatoire des séries formelles, *Adv. in Math.* **42** (1981), 1-82.

8. A. Lindenbaum and A. Tarski, Communication sur les recherches de la theorie des ensembles, in "Alfred Tarski: collected papers," Birkhauser, Boston, 1986.

9. S. Mac Lane, "Categories for the Working Mathematician," Springer-Verlag, New York, 1971.

10. S. Mac Lane and I. Moerdijk, "Sheaves in Geometry and Logic, A First Introduction to Topos Theory," Springer-Verlag, New York, 1992.

11. R. Stanley, "Enumerative Combinatorics I," Wadsworth, Monterey, 1986.

12. H. Wilf, Sieve-equivalence in generalized partition theory, *J. Comb. Th. Ser. A* **34** (1983), 80-89.