Finite Preimages Under the Natural Map from $\beta(\mathbb{N} \times \mathbb{N})$ to $\beta\mathbb{N} \times \beta\mathbb{N}$

Andreas Blass*
Mathematics Department
University of Michigan
Ann Arbor, MI 48109–1109, U.S.A.
ablass@umich.edu

Gugu Moche
Mathematics Department
Howard University
Washington, DC 20059
waazania@aol.com

January 22, 2003

Abstract

Let $\widetilde{\iota}: \beta(\mathbb{N} \times \mathbb{N}) \to \beta\mathbb{N} \times \beta\mathbb{N}$ be the continuous extension of the identity map of $\mathbb{N} \times \mathbb{N}$. We provide elementary proofs of several sharp results about the possible sizes of preimages of points $(p, q) \in \mathbb{N}^* \times \mathbb{N}^*$. Among these are:

1. If $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 2$ then both p and q are P-points, but assuming the existence of at least 3 non-isomorphic, selective ultrafilters on $\mathbb{N}$, there are p and q such that $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 3$ but p is not a P-point.
2. If $|\widetilde{\iota}^{-1}[\{(p,q)\}]| \leq 5$ then at least one of p and q is a P-point, but assuming the existence of at least 4 non-isomorphic selective

*Partially supported by NSF grant DMS–0070723

ultrafilters on $\mathbb{N}$, there are p and q such that neither of them is a P-point and $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 6$.

3. If $|\widetilde{\iota}^{-1}[\{(p,p)\}]| \leq 8$ then p is a P-point, but assuming the existence of infinitely many non-isomorphic, selective ultrafilters on $\mathbb{N}$, there is a non-P-point p such that $|\widetilde{\iota}^{-1}[\{(p,p)\}]| = 9$.

1 Introduction

We shall be concerned with the Stone-Čech compactifications βX of countable sets X, where X is always taken to have the discrete topology. In particular, we consider $\beta\mathbb{N}$ and $\beta(\mathbb{N}^k)$ where $\mathbb{N}$ is the set of natural numbers. We use the customary identification of points of βX with ultrafilters on X, the points of X being identified with principal ultrafilters.

Any function f from X into any compact Hausdorff space Z extends uniquely to a continuous map $\widetilde{f} : \beta X \to Z$. In particular, a function f from X to any set Y can be regarded as mapping into βY, and then it has a continuous extension $\widetilde{f} : \beta X \to \beta Y$. Another special case, which will play a central role in this paper, is given by the identity map of $\mathbb{N} \times \mathbb{N}$, viewed as a function $\iota : \mathbb{N}\times\mathbb{N} \to \beta\mathbb{N}\times\beta\mathbb{N}$, and its extension $\widetilde{\iota} : \beta(\mathbb{N}\times\mathbb{N}) \to \beta\mathbb{N}\times\beta\mathbb{N}$. We shall study points $(p,q) \in \beta\mathbb{N} \times \beta\mathbb{N}$ whose preimage $\widetilde{\iota}^{-1}[\{(p,q)\}]$ is finite; we relate this finiteness property to other combinatorial properties of ultrafilters, particularly the concepts of P-points and selective ultrafilters.

In the 1970's, the possible cardinalities of $\widetilde{\iota}^{-1}[\{(p,q)\}]$ were investigated by the first author in [1, 2], by Neil Hindman in [9], and by Maryvonne Daguenet in [6].

In 2000 and 2001, the second author was investigating the question whether the restriction of $\widetilde{\iota}$ to the smallest ideal (see [10]) of the compact right-topological semigroup $(\beta(\mathbb{N} \times \mathbb{N}), +)$ might be finite-to-one or even two-to-one. An affirmative answer would have been very useful in describing that smallest ideal algebraically; unfortunately, the answer turned out to be negative [12].

In the course of her investigation, the second author was made aware of the substantial results that were known to the experts about the possible sizes of $\widetilde{\iota}^{-1}[\{(p,q)\}]$. Unfortunately, many of these results were (1) published in not widely available sources like [1, 6], or (2) published with model-theoretic proofs that require substantial training in mathematical logic to follow, or (3) not published at all. Specifically, the existence parts of statements (1) and (2) in the abstract appear not to have been published previously, and the first part of statement (3) has apparently been published only with a model-theoretic proof.

Because these results are very sharp and of interest from both a combinatorial and a topological point of view, we feel that it is worthwhile to make their statements and proofs more accessible. In this paper, we present proofs of the known results about the cardinalities of preimages $\widetilde{\iota}^{-1}[\{(p,q)\}]$, with the exception of those results for which elementary proofs are already in the widely available literature.

In Section 2 we review the terminology, notation, and basic results that we shall use. In Section 3, we show that the existence of finite preimages $\widetilde{\iota}^{-1}[\{(p,q)\}]$ (except in the trivial case where p or q is principal) requires the existence of P-point ultrafilters and therefore cannot be proved from the usual (ZFC) axioms of set theory. In Section 4, we give more specific results, relating the cardinality of $\widetilde{\iota}^{-1}[\{(p,q)\}]$ to the P-point property of p or q or both. These results do not require any assumptions beyond ZFC, but they are vacuous if there are no P-points. Finally, in Section 5, we show that the results of Section 4 are optimal provided there are infinitely many non-isomorphic selective ultrafilters. This proviso is satisfied if the continuum hypothesis holds, and also under weaker hypotheses such as Martin's Axiom, but it is not provable in ZFC.

Acknowledgment

We thank Neil Hindman, the second author's thesis advisor, for his guidance of her research and for suggesting that an elementary presentation such as the present paper would be a worthwhile undertaking.

2 Definitions and Preliminary Results

In this section, we review some well-known definitions and facts that we shall need later. A general reference for most of this material is [5].

2.1 Images of ultrafilters

For any set X, regarded as a discrete space, the Stone-Čech compactification βX has a basis of open sets consisting of the closures $\overline{A}$ of subsets of X.

Under the identification of points of βX with ultrafilters on X, $\overline{A}$ consists of those ultrafilters that contain A. We use the customary notation X^* for the space $\beta X - X$ of non-principal ultrafilters on X.

We shall need some elementary facts about the continuous extensions $\widetilde{f} : \beta X \to \beta Y$ of maps $f : X \to Y$. If $f : X \to Y$ and $g : Y \to Z$, then $\widetilde{g \circ f} = \widetilde{g} \circ \widetilde{f}$ because both sides of this equation are continuous extensions of $g \circ f$. This fact and the trivial observation that identity maps id_X have $\widetilde{\mathrm{id}_X} = \mathrm{id}_{\beta X}$ together imply that, if f is a bijection (with inverse g), then $\widetilde{f}$ is a bijection (with inverse $\widetilde{g}$), in fact a homeomorphism.

In terms of ultrafilters, the extension $\widetilde{f}$ of $f : X \to Y$ can be defined by

$$\widetilde{f}(p) = \{B \subseteq Y : f^{-1}[B] \in p\};$$

this ultrafilter is generated by the sets $f[A]$ for $A \in p$.

We use the notation π_k for the k^{th} projection map of a product,

$$\pi_k : X_1 \times \cdots \times X_n \to X_k : (x_1, \ldots, x_n) \mapsto x_k.$$

This could lead to ambiguity when a product can be broken into factors in more than one way. We resolve the only ambiguous cases arising in this paper by agreeing that the projections π_k of $\mathbb{N}^n$ are its n projections to $\mathbb{N}$. When we need projections to subproducts, we'll introduce a separate notation for them.

Definition 1 The *Rudin-Keisler ordering* of ultrafilters is the pre-order defined by setting $p \leq_{RK} q$ (where p and q are ultrafilters on X and Y respectively) if $p = \widetilde{f}(q)$ for some $f : Y \to X$. p and q are *isomorphic*, written $p \cong q$, if $p = \widetilde{f}(q)$ for some f that is one-to-one on a set in q.

Both $\leq_{RK}$ and $\cong$ are reflexive and transitive trivially. The latter is also symmetric, since if f witnesses that $p \cong q$ and f is one-to-one on $A \in q$, then any $g : X \to Y$ that agrees with $(f \restriction A)^{-1}$ on $f[A]$ witnesses that $q \cong p$. We have $p \cong q$ if and only if $p \leq_{RK} q \leq_{RK} p$; here "only if" is obvious and "if" is an immediate consequence of the following well-known result, for whose proof we refer to [5, Theorem 9.2(a)].

Lemma 2 *If* $\widetilde{f}(p) = p$ *then* $\{x : f(x) = x\} \in p$.

Lemma 3 *If* p *and* q *are non-principal ultrafilters on countably infinite sets* X *and* Y, *then* $p \cong q$ *if and only if* $p = \widetilde{f}(q)$ *for some bijection* $f : Y \to X$.

Proof The "if" part is immediate from the definition of $\cong$. For the "only if" part, assume that $p \cong q$ and fix a $g : Y \to X$ such that $p = \widetilde{g}(q)$ and g is one-to-one on a set $A \in q$. Split A into two infinite pieces; exactly one of them, say B, will be in q. Then $Y - B$ and $X - g[B]$ are countably infinite, because they include the infinite set $A - B$ and its one-to-one image $g[A - B]$, respectively. Let $f : Y \to X$ agree with g on B and with an arbitrary bijection between the countably infinite sets $Y - B$ and $X - g[B]$. Then f is a bijection from Y to X. Since it agrees with g on a set $B \in q$ we have $\widetilde{f}(q) = \widetilde{g}(q) = p$. $\square$

2.2 The natural map from $\beta(\mathbb{N} \times \mathbb{N})$ to $\beta\mathbb{N} \times \beta\mathbb{N}$

The map $\widetilde{\iota}$ does not fit into the general discussion of $\widetilde{f}$ above, since the map $\iota : \mathbb{N} \times \mathbb{N} \to \beta\mathbb{N} \times \beta\mathbb{N}$ is not obtained by regarding a map $X \to Y$ as a map $X \to \beta Y$. But the two projections of ι are obtained in this way, from the two projections of $\mathbb{N} \times \mathbb{N}$. Thus, $\widetilde{\iota}(s) = (\widetilde{\pi}_1(s), \widetilde{\pi}_2(s))$ for all ultrafilters s on $\mathbb{N} \times \mathbb{N}$. Equivalently, we have $\widetilde{\iota}(s) = (p, q)$ if and only if s contains all sets of the form $A \times B$ with $A \in p$ and $B \in q$.

It is easy to check that, if either p or q is principal, then there is only one such s; the sets $A \times B$ as above generate an ultrafilter. Thus, we shall be interested in the cardinality of $\widetilde{\iota}^{-1}[\{(p, q)\}]$ only when both p and q are non-principal.

The following is an easy consequence of Lemma 3.

Corollary 4 *If $q \cong r$ are non-principal ultrafilters on $\mathbb{N}$, then $|\widetilde{\iota}^{-1}[\{(p, q)\}]| = |\widetilde{\iota}^{-1}[\{(p, r)\}]|$.*

Proof By Lemma 3, fix a permutation f of $\mathbb{N}$ such that $\widetilde{f}(r) = q$. Then $f'(x, y) = (x, f(y))$ defines a permutation of $\mathbb{N}^2$, inducing a self-homeomorphism $\widetilde{f'}$ of $\beta(\mathbb{N}^2)$. The equations $\widetilde{\pi}_1 \circ \widetilde{f'} = \widetilde{\pi}_1$ and $\widetilde{\pi}_2 \circ \widetilde{f'} = \widetilde{f} \circ \widetilde{\pi}_2$ (which follow from the corresponding equations without the tildes) immediately imply that $\widetilde{f'}$ maps $\widetilde{\iota}^{-1}[\{(p, r)\}]$ to $\widetilde{\iota}^{-1}[\{(p, q)\}]$. Applying the same argument starting with f^{-1}, we find that $\widetilde{f'}$ is a bijection from $\widetilde{\iota}^{-1}[\{(p, r)\}]$ to $\widetilde{\iota}^{-1}[\{(p, q)\}]$. $\square$

Lemma 5 *Let p and q be non-principal ultrafilters on $\mathbb{N}$, and let k be a natural number. The following are equivalent.*

1. *$\widetilde{\iota}^{-1}[\{(p, q)\}]$ has at most k members.*

2. *For any natural number r and any map $f : \mathbb{N}\times\mathbb{N} \to \{1, 2, \ldots, r\}$, there exist $A \in p$ and $B \in q$ such that $f[A \times B]$ has at most k members.*

3. *If $C_1, \ldots, C_{k+1}$ are $k + 1$ pairwise disjoint subsets of $\mathbb{N} \times \mathbb{N}$, then one of them is disjoint from some set of the form $A \times B$ with $A \in p$ and $B \in q$.*

Proof If (1) fails and we have $k+1$ distinct ultrafilters $s_i \in \widetilde{\iota}^{-1}[\{(p, q)\}]$ then disjoint basic neighborhoods $\overline{C_i}$ of these ultrafilters give sets C_i witnessing that (3) fails.

If sets C_i form a counterexample to (3), then any function that has value i on C_i for each i will witness that (2) fails.

Finally, if f is a counterexample to (2) then for at least $k + 1$ values of i the family $\{f^{-1}[\{i\}]\} \cup \{A \times B : A \in p \text{ and } B \in q\}$ has the finite intersection property and is therefore included in an ultrafilter s_i. These s_i are at least $k + 1$ distinct elements of $\widetilde{\iota}^{-1}[\{(p, q)\}]$, so (1) fails. □

Corollary 6 *Let p and q be non-principal ultrafilters on $\mathbb{N}$. Then $|\widetilde{\iota}^{-1}[\{(p, q)\}]| \geq 2$ and $|\widetilde{\iota}^{-1}[\{(p, p)\}]| \geq 3$.*

Some sets involved in the proof of this corollary will be needed again later, so we provide a permanent notation for them (and for a useful map) before proving the corollary.

Definition 7 We use U (for "up"), D (for "down"), and Δ (for "diagonal) to denote the following subsets of $\mathbb{N}^2$.

$$\begin{aligned} U &= \{(x, y) \in \mathbb{N} \times \mathbb{N} : x < y\}, \\ D &= \{(x, y) \in \mathbb{N} \times \mathbb{N} : x > y\}, \text{ and} \\ \Delta &= \{(x, y) \in \mathbb{N} \times \mathbb{N} : x = y\}. \end{aligned}$$

We write $\tau : \mathbb{N}^2 \to \mathbb{N}^2$ for the "twist" map interchanging the two coordinates, $\tau(x, y) = (y, x)$.

Proof of Corollary Assertion (3) of the lemma is false for $k = 1$, since we can take $C_1 = U$ and $C_2 = D$, both of which meet every set of the form $A \times B$ with A and B infinite. If $q = p$, then as the diagonal Δ also meets all sets of the form $A \times A$ with A nonempty, we have a violation of (3) for $k = 2$. □

Lemma 8 *If $|\widetilde{\iota}^{-1}[\{(p,p)\}]|$ is finite then it is odd. In fact, it is $2k+1$ where $k = |\overline{U} \cap \widetilde{\iota}^{-1}[\{(p,p)\}]|$.*

Proof As U, D, and Δ partition $\mathbb{N}^2$, their closures partition $\beta(\mathbb{N}^2)$, so we can count the ultrafilters in $\widetilde{\iota}^{-1}[\{(p,p)\}]$ by counting separately those in $\overline{U}$, $\overline{D}$, and $\overline{\Delta}$. Since $\tau \circ \tau = \mathrm{id}$, $\pi_1 \circ \tau = \pi_2$, and $\pi_2 \circ \tau = \pi_1$, the corresponding equations hold for the continuous extensions $\widetilde{\tau}$ and $\widetilde{\pi}_i$. It follows that $\widetilde{\tau}$ is an involution (i.e., a bijection that is its own inverse) of $\beta(\mathbb{N}^2)$, sending $\widetilde{\iota}^{-1}[\{(p,p)\}]$ onto itself (more generally, sending $\widetilde{\iota}^{-1}[\{(p,q)\}]$ onto $\widetilde{\iota}^{-1}[\{(q,p)\}]$), and interchanging $\overline{U}$ with $\overline{D}$. Thus, if $\widetilde{\iota}^{-1}[\{(p,p)\}]$ has k elements in $\overline{U}$, then it has another k elements in $\overline{D}$. It also has exactly one member in $\overline{\Delta}$, namely the image of p under the map $\mathbb{N} \to \Delta : n \mapsto (n,n)$. So the total number of ultrafilters in $\widetilde{\iota}^{-1}[\{(p,p)\}]$ is $2k+1$. □

We omit the proof of the following lemma as all but the last part is analogous to that of Lemma 5. The last part is trivial once one notices that U is essentially the same as the set $[\mathbb{N}]^2$ of two-element subsets of $\mathbb{N}$, with $U \cap (A \times A)$ corresponding to $[A]^2$.

Lemma 9 *Let p be a non-principal ultrafilter on $\mathbb{N}$, and let k be a natural number. The following are equivalent.*

1. *$\widetilde{\iota}^{-1}[\{(p,p)\}]$ has at most $2k+1$ members.*
2. *For any natural number r and any map $f : U \to \{1, 2, \ldots, r\}$, there exists $A \in p$ such that $f[(A \times A) \cap U]$ has at most k members.*
3. *If $C_1, C_2 \ldots, C_{k+1}$ are $k+1$ pairwise disjoint subsets of U, then one of them is disjoint from some set of the form $A \times A$ with $A \in p$.*
4. *If $[\mathbb{N}]^2$ is partitioned into finitely many pieces, then there is $A \in p$ such that $[A]^2$ meets at most k of the pieces.*

2.3 Special ultrafilters

Definition 10 A *P-point* is a non-principal ultrafilter p on $\mathbb{N}$ such that every function $\mathbb{N} \to \mathbb{N}$ becomes finite-to-one or constant when restricted to some set in p. A non-principal ultrafilter p on $\mathbb{N}$ is *selective* if every function $\mathbb{N} \to \mathbb{N}$ becomes one-to-one or constant when restricted to some set in p.

Obviously, every selective ultrafilter is a P-point. An equivalent characterization of P-points is that, for every sequence $\langle A_n \rangle_{n\in\mathbb{N}}$ of members of p, there is $B \in p$ that is almost included in each A_n, meaning that $B - A_n$ is finite. An ultrafilter is selective if and only if it is minimal with respect to the Rudin-Keisler ordering among all non-principal ultrafilters.

It is known that the continuum hypothesis implies the existence of infinitely many (in fact $2^{\mathfrak{c}}$ where $\mathfrak{c}$ is the cardinality of the continuum) selective ultrafilters and also of P-points that are not selective. (In fact, weaker assumptions than CH suffice for this, for example Martin's Axiom [3] or $\mathfrak{p} = \mathfrak{c}$ [1].) But the existence of P-points cannot be proved in ZFC alone, without special assumptions [15]. It is also consistent with ZFC that there are P-points but no selective ultrafilters; specifically, it shown in [11] that there are no selective ultrafilters in the so-called random real model, and it is shown in [4] that there are P-points in this model.

Lemma 11 *If a non-principal ultrafilter p on $\mathbb{N}$ is not a P-point, then there is a function $f : \mathbb{N} \to \mathbb{N}$ such that $f(x) \leq x$ for all $x \in \mathbb{N}$ and such that f is neither finite-to-one nor constant on any set in p.*

Proof The definition of "P-point" provides an $f_0 : \mathbb{N} \to \mathbb{N}$ that is neither finite-to-one nor constant on any set in p. Define $f(x) = \min\{f_0(x), x\}$ and observe that this f is as required in the lemma. Indeed, if f were constant with value n on a set $A \in p$ then f_0 would be constant on $A - \{n\} \in p$, and if f were finite-to-one on $A \in p$, then f_0 would also be finite-to-one on the same A, taking each value at most one more time than f did. □

2.4 Products of ultrafilters

We shall need the following construction of a product of two ultrafilters.

Definition 12 Let p and q be ultrafilters on X and Y, respectively. Then their *tensor product* is the ultrafilter

$$p \otimes q = \{A \subseteq X \times Y : \{x \in X : A_x \in q\} \in p\},$$

where A_x means the "slice" $A_x = \{y \in Y : (x, y) \in A\}$.

An equivalent form of the definition that may be easier to remember is that a set $A \subseteq X \times Y$ contains almost all pairs (x, y) with respect to $p \otimes q$

if and only if for almost all x with respect to p, for almost all y with respect to q, $(x, y) \in A$.

It is easy to verify that $p \otimes q$ is an ultrafilter and that its images under the projection maps are $\widetilde{\pi}_1(p \otimes q) = p$ and $\widetilde{\pi}_2(p \otimes q) = q$. Thus, if $p, q \in \beta\mathbb{N}$ then $p \otimes q \in \widetilde{\iota}^{-1}[\{(p, q)\}]$.

We shall need a characterization of tensor products due to Puritz [13]; Puritz's proof is model-theoretic, so we give a combinatorial one.

Proposition 13 *Let X and Y be countably infinite sets and let $p \in X^*$ and $q \in Y^*$. Then $p \otimes q$ is the unique ultrafilter $r \in \beta(X \times Y)$ such that*

1. *$\widetilde{\pi}_1(r) = p$,*

2. *$\widetilde{\pi}_2(r) = q$, and*

3. *For all functions $f : X \to \mathbb{N}$ and $g : Y \to \mathbb{N}$, either g is constant on a set in q or $\{(x, y) : f(x) < g(y)\} \in r$.*

Proof We have already observed that $p \otimes q$ has properties 1 and 2. To see that it also has property 3, observe that, if g is not constant on any set in q, then it is also not bounded above on any set in q (as q is an ultrafilter), and so the sets $\{y : f(x) < g(y)\}$ are in q for all x. As these sets are the slices of $\{(x, y) : f(x) < g(y)\}$, this set is in $p \otimes q$.

Now suppose r has the three properties in the proposition, and suppose, toward a contradiction, that $r \neq p \otimes q$. So there is a set $A \in r$ that does not belong to $p \otimes q$, which means that the set $B = \{x \in X : A_x \notin q\}$ is in p. By property 1 of r, it follows that r contains ${\pi_1}^{-1}(B) = B \times Y$. So $A' = A \cap (B \times Y)$ is in r, and being a subset of A it is certainly not in $p \otimes q$. No slices ${A'}_x$ of A' are in q, for wherever a slice A_x was in q (i.e., $x \notin B$), the corresponding slice of A' is empty.

Fix arbitrary bijections $f : X \to \mathbb{N}$ and $h : Y \to \mathbb{N}$, and let

$$A'' = A' \cup \{(x, y) \in X \times Y : f(x) = h(y)\}.$$

Being a superset of A', this A'' is in r. Its slices differ from those of A' by singletons (as f and h are bijections), so they are not in q.

Replacing the original A by A'', we may assume without loss of generality that $A \in r$, that $A_x \notin q$, and that $h^{-1}(f(x)) \in A_x$ for all $x \in X$. This last property allows us to define a function $g : Y \to \mathbb{N}$ by

$$g(y) = \text{the smallest } n \in \mathbb{N} \text{ such that } y \in A_{f^{-1}(n)}.$$

Indeed, taking $x = f^{-1}(h(y))$ in the statement $h^{-1}(f(x)) \in A_x$ above, we get $y \in A_{f^{-1}(h(y))}$; thus, $h(y)$ is an n of the sort required in the definition of $g(y)$. We conclude that $g(y)$ exists (and in fact $g(y) \leq h(y)$).

Consider any pair $(x, y) \in A$. Then $f(x)$ is an n of the sort required in the definition of $g(y)$, i.e., $y \in A_x = A_{f^{-1}(f(x))}$. Therefore $g(y) \leq f(x)$. Since $A \in r$, we do not have the second alternative in property 3 of r. So we must have the first alternative, namely that g is constant, say with value n, on some set $C \in q$. Then by definition of g we have that $C \subseteq A_{f^{-1}(n)}$. But this is absurd, as $C \in q$ and $A_{f^{-1}(n)} \notin q$. This contradiction completes the proof of Puritz's theorem. $\square$

3 Finite Preimages Require P-Points

The purpose of this section is to show that finiteness of $\widetilde{\iota}^{-1}[\{(p, q)\}]$ requires the existence of P-points.

Theorem 14 *Assume that p and q are non-principal ultrafilters on $\mathbb{N}$ and that $\widetilde{\iota}^{-1}[\{(p, q)\}]$ is finite. Then there exists a P-point. In fact, there must be a P-point $\leq_{RK} p$ and a P-point $\leq_{RK} q$.*

The essence of the proof of this theorem is given in [6, Theorem 11], though the statement of the theorem is rather different.

Proof Since the roles of p and q are symmetrical, it suffices to deduce a contradiction from the assumption that there is no P-point $\leq_{RK} p$.

We begin by defining a sequence of functions $\langle f_n \rangle_{n \in \mathbb{N}}$ where $f_n : \mathbb{N} \to \mathbb{N}$, $f_n(x) \leq x$ for all $x \in \mathbb{N}$, and f_n is neither finite-to-one nor constant on any set in $\widetilde{h_n}(p)$, where $h_1 = \mathrm{id}_{\mathbb{N}}$ and $h_{n+1} = f_n \circ h_n$. (In other words, $h_n = f_{n-1} \circ \cdots \circ f_1$.) The definition of f_n is by induction on n, so assume that we have already defined f_k for all $k < n$; thus h_n is already defined. Also assume as an induction hypothesis that $\widetilde{h_n}(p)$ is non-principal. (Note that this is true when $n = 1$.) Since we are assuming, toward a contradiction, that there is no P-point $\leq_{RK} p$, we know in particular that $\widetilde{h}_n(p)$ is not a P-point. By Lemma 11, there is a function f_n that satisfies $f_n(x) \leq x$ for all $x \in \mathbb{N}$ and that is neither finite-to-one nor constant on any set in $\widetilde{h_n}(p)$. In particular, as it isn't constant, the image ultrafilter $\widetilde{f}_n(\widetilde{h_n}(p)) = \widetilde{h_{n+1}}(p)$ is non-principal, so our induction hypothesis is preserved. This completes the construction of the f_n's and thus of the h_n's.

Notice that, for all n and all x, $h_{n+1}(x) = f_n(h_n(x)) \leq h_n(x)$.

Define, for each $n \in \mathbb{N}$,

$$C_n = \{(x, y) \in \mathbb{N}^2 : h_{n+1}(x) < y < h_n(x)\},$$

and notice that these sets are pairwise disjoint. We intend to show that each C_n meets every set of the form $A \times B$ for $A \in p$ and $B \in q$. By Lemma 5, this will imply that $\widetilde{\iota}^{-1}[\{(p, q)\}]$ is infinite, contrary to the hypothesis of the theorem, so the proof will be complete.

So fix any $n \in \mathbb{N}$, $A \in p$, and $B \in q$. Being in the non-principal ultrafilter $\widetilde{h_n}(p)$, the set $h_n[A]$ is infinite, and f_n is, by construction, not finite-to-one on this set. So we can choose a number z that is $f_n(w)$ for infinitely many $w \in h_n[A]$. Since B is infinite, we can choose $y \in B$ with $y > z$. Then we can choose $w \in h_n[A]$ with $f_n(w) = z$ and $w > y$, and we can choose $x \in A$ with $h_n(x) = w$. Then $h_{n+1}(x) = f_n(w) = z < y < w = h_n(x)$. So $(x, y) \in (A \times B) \cap C_n$, as required. □

Corollary 15 *It is consistent relative to ZFC that $\widetilde{\iota}^{-1}[\{(p, q)\}]$ is infinite for all non-principal ultrafilters p, q on $\mathbb{N}$.*

Proof Combine Theorem 14 with the theorem of Shelah [15] that it is consistent relative to ZFC that there are no P-points. □

Remark 16 The idea in the proof of Theorem 14 can be modified to give an alternative proof of a theorem of Hindman [9, Theorem 3.1], namely that there exists an ultrafilter p on $\mathbb{N}$ such that $\widetilde{\iota}^{-1}[\{(p, q)\}]$ is infinite for all $q \in \mathbb{N}^*$. We briefly sketch this proof.

Fix a function $f : \mathbb{N} \to \mathbb{N}$ such that $f(x) \leq x$ and $f^{-1}[\{x\}]$ is infinite for all $x \in \mathbb{N}$. The idea is to let this f play the roles of all the f_n's in the proof of Theorem 14, so that h_n is f^{n-1}. The last part of the proof of the theorem shows that $\widetilde{\iota}^{-1}[\{(p, q)\}]$ is infinite provided f is not constant or finite-to-one on any set in $\widetilde{h_n}(p)$. So our objective is to build an ultrafilter p containing no sets of the forms ${h_{n+1}}^{-1}\{x\}$ (for any $x \in \mathbb{N}$) and $\{x : h_n(x) \leq g(h_{n+1}(x))\}$ (for any $g : \mathbb{N} \to \mathbb{N}$); indeed the sets of the first (resp. second) form and their subsets are exactly those A such that f_n is constant (resp. finite-to-one) on $h_n[A]$.

Thus, the required p will exist provided $\mathbb{N}$ cannot be covered by finitely many sets of the two prohibited forms. But it is easy to check that this proviso is satisfied, because all $f^{-1}[\{x\}]$ are infinite.

4 The Smallest Preimages

In this section, we obtain results more precise than Theorem 14. Assuming certain specific bounds on $|\widetilde{\iota}^{-1}[\{(p,q)\}]|$ (rather than merely assuming that it is finite), we obtain P-points among the ultrafilters p, q themselves (rather than merely below them in the Rudin-Keisler order). In the next section, we shall see that the bounds used here are sharp.

Recall that, for non-principal ultrafilters on $\mathbb{N}$, the smallest possible size for $\widetilde{\iota}^{-1}[\{(p,q)\}]$ is 2. The following theorem tells us exactly when this minimum is attained. It was proved in [6]; the "P-point" part of the "only if" half of this theorem was proved earlier in both [8] and [1].

Theorem 17 *For non-principal ultrafilters p, q on $\mathbb{N}$, $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 2$ if and only if p and q are P-points and there is no non-principal r that is $\leq_{RK}$ both of them.*

Proof Assume that $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 2$. We show first that p is a P-point. So let $f : \mathbb{N} \to \mathbb{N}$ be arbitrary; we seek a set in p on which f is finite-to-one or constant. Consider the following three disjoint subsets of $\mathbb{N}^2$ (of which the first is as in Definition 7 and the other two are parts of the D from that definition):

$$\begin{aligned} U &= \{(x,y) : x < y\} \\ D_1 &= \{(x,y) : x > y \text{ and } y < f(x)\} \\ D_2 &= \{(x,y) : x > y \text{ and } y > f(x)\}. \end{aligned}$$

By Lemma 5, we have some $A \in p$ and $B \in q$ such that $A \times B$ meets at most two of U, D_1, D_2.

Take any $x \in A$ and then, as B is infinite, find $y \in B$ such that $y > x$. That gives a pair in $(A \times B) \cap U$. So $A \times B$ is disjoint from one of D_1 and D_2.

Suppose $(A \times B) \cap D_1 = \varnothing$. Take any $y \in B$ and observe that all $x \in A - [1,y]$ must satisfy $f(x) \leq y$ (lest (x,y) be in D_1). So f is bounded on a set $A - [1,y]$ in p; since p is an ultrafilter it follows that f is constant on a set in p.

There remains the case that $(A \times B) \cap D_2 = \varnothing$. In this case, given any $z \in \mathbb{N}$, we can find $y \in B$ such that $y > z$, and then we can infer that every $x \in A \cap f^{-1}[\{z\}]$ must be $\leq y$ (lest (x,y) be in D_2 because $x > y > z = f(x)$).

So f is finite-to-one on $A \in p$. This completes the proof that p is a P-point. The proof for q is symmetrical.

It remains to show that there is no non-principal $r \leq_{RK} p, q$. Suppose, toward a contradiction, that we had $r = \widetilde{f}(p) = \widetilde{g}(q)$. We obtain a contradiction via Lemma 5, by exhibiting three disjoint subsets of $\mathbb{N}^2$, each of which meets $A \times B$ for all $A \in p$ and $B \in q$. The three sets are

$$\begin{aligned} C_1 &= \{(x,y) \in U : f(x) \neq g(y)\} \\ C_2 &= \{(x,y) \in D : f(x) \neq g(y)\} \\ C_3 &= \{(x,y) \in \mathbb{N}^2 : f(x) = g(y)\}. \end{aligned}$$

To show that these do the job, consider any $A \in p$ and $B \in q$. Pick any $x \in A$. Since r is non-principal, g is not constant on any set in q. In particular, g does not map all elements of $B - [1, x]$ to $f(x)$. So pick $y \in B - [1, x]$ with $g(y) \neq f(x)$ and observe that $(x, y) \in C_1$. Symmetrically, $A \times B$ meets C_2. Finally, since r contains both $f[A]$ and $g[B]$, we can choose an element $z \in f[A] \cap g[B]$, and then we can choose $x \in A$ and $y \in B$ such that $f(x) = z = g(y)$. Thus, $A \times B$ meets C_3 also.

This completes the proof of the "only if" half of the theorem. The "if" half will use the following lemma, which will be needed again later.

Lemma 18 *Suppose r is a non-principal ultrafilter on a countable set X, and suppose $f_1, f_2 : X \to \mathbb{N}$ are functions such that*

$$\{x \in X : f_1(x) \leq h_1(f_2(x))\} \in r \quad \text{and} \quad \{x \in X : f_2(x) \leq h_2(f_1(x))\} \in r$$

for some functions $h_1, h_2 : \mathbb{N} \to \mathbb{N}$. Then there exist finite-to-one functions $g_1, g_2 : \mathbb{N} \to \mathbb{N}$ such that

$$\{x \in X : g_1(f_1(x)) = g_2(f_2(x))\} \in r.$$

In other words, if each of f_1 and f_2 is majorized, on a set in r, by a function of the other, then they become equal on a set in r when composed with suitable finite-to-one maps.

Proof of Lemma We may assume h_1 and h_2 are the same function h; just take the pointwise maximum of the two functions. Inductively choose an increasing sequence $1 = a_1 < a_2 < a_3 < \dots$ of natural numbers such that $h(k) < a_{n+1}$ for all $k \leq a_n$. Since almost all (with respect to r) x's satisfy $f_1(x) \leq h(f_2(x))$, they have the property that, if $f_1(x) \in [a_n, a_{n+1})$ and

$f_2(x) \in [a_m, a_{m+1})$, then $m \leq n+1$; symmetrically, we also have $n \leq m+1$ for almost all x.

Let us define $g : \mathbb{N} \to \mathbb{N}$ as the step function whose value is n on the interval $[a_n, a_{n+1})$. Then we have shown that

$$\{x \in X : g(f_1(x)) = g(f_2(x)) \text{ or } g(f_2(x)) + 1 \text{ or } g(f_2(x)) - 1\} \in r.$$

As r is an ultrafilter, it must contain one of the three sets

$$\begin{aligned} &\{x : g(f_1(x)) = g(f_2(x))\}, \\ &\{x : g(f_1(x)) = g(f_2(x)) + 1\}, \\ &\{x : g(f_1(x)) = g(f_2(x)) - 1\}. \end{aligned}$$

If it contains the first, then we can satisfy the conclusion of the lemma by taking $g_1 = g_2 = g$. If it contains the second, then take $g_1 = g$ and $g_2 = g+1$. And if it satisfies the third, then take $g_1 = g+1$ and $g_2 = g$. $\square$

We now resume the proof of Theorem 17. Assume p and q are P-points in $\mathbb{N}^*$ and there is no non-principal $r \leq_{RK} p, q$. Consider an arbitrary $s \in \widetilde{\iota}^{-1}[\{(p,q)\}]$.

Temporarily suppose $U \in s$. We intend to show under this assumption that $s = p \otimes q$, by applying Theorem 13. The first two hypotheses of that theorem are satisfied by assumption, so consider arbitrary functions $f, g : \mathbb{N} \to \mathbb{N}$. We must show that either g is constant on a set in q or $f(x) < g(y)$ on a set in s. Suppose, toward a contradiction, that neither of these alternatives holds. As q is a P-point, the failure of the first alternative means that g is finite-to-one on some set $B \in q$. So we can define a function $h : \mathbb{N} \to \mathbb{N}$ such that $h(n) > f(x)$ whenever $x < y$, $y \in B$, and $g(y) = n$. The reason is that, for any n, there are only finitely many $y \in B$ with $g(y) = n$, and for each of these y's there are only finitely many $x < y$. Having defined h in this way, we see that $f(x) < h(g(y))$ for all $(x, y) \in U \cap (\mathbb{N} \times B)$. Note that this set is in s, as we have $U \in s$ and $B \in q = \widetilde{\pi}_2(s)$.

Furthermore, the failure of the second alternative gives us that $g(y) \leq f(x)$ on a set in s. Thus, Lemma 18 is applicable with s in the role of r, with $f \circ \pi_1$ and $g \circ \pi_2$ in the roles of the f_i, and with h and $\mathrm{id}_{\mathbb{N}}$ in the roles of the h_i. We conclude that there are finite-to-one functions g_i such that $g_1 \circ f \circ \pi_1 = g_2 \circ g \circ \pi_2$ on a set in s. Since $\widetilde{\pi}_1(s) = p$ and $\widetilde{\pi}_2(s) = q$, it follows that $\widetilde{g_1 \circ f}(p) = \widetilde{g_2 \circ g}(q)$, so we have found an ultrafilter that is $\leq_{RK} p, q$. By hypothesis, this ultrafilter must be principal. But in fact, this ultrafilter

$\widetilde{g_2 \circ g}(q)$ cannot be principal because $g_2 \circ g$ is finite-to-one on a set in q and therefore cannot be constant on a set in q.

This contradiction completes the proof that $s = p \otimes q$ under the temporary assumption that $U \in s$. Since the assumptions are symmetrical between p and q, we also have that $q \otimes p$ is the only ultrafilter in $\widetilde{\iota}^{-1}[\{(q,p)\}]$ that contains U; therefore $\widetilde{\tau}(q \otimes p)$ is the only ultrafilter in $\widetilde{\iota}^{-1}[\{(p,q)\}]$ that contains $\tau[U] = D$. No ultrafilter in $\widetilde{\iota}^{-1}[\{(p,q)\}]$ can contain Δ, for such an ultrafilter would be isomorphic, via the projection maps, to both p and q, which clearly contradicts the hypotheses of the theorem.

Therefore, there are exactly two ultrafilters in $\widetilde{\iota}^{-1}[\{(p,q)\}]$, namely $p \otimes q$ and $\widetilde{\tau}(q \otimes p)$. $\square$

We shall see in the next section that it is consistent to have $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 3$ while only one of p, q is a P-point.

Corollary 19 *If p and q are non-isomorphic, selective ultrafilters on $\mathbb{N}$, then $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 2$.*

Proof Selective ultrafilters are P-points, so it suffices to check that there is no non-principal $r \leq_{RK} p, q$. By selectivity, any non-principal ultrafilter $\leq_{RK} p$ is in fact $\cong p$, and similarly for q. So any $r \leq_{RK} p, q$ would be isomorphic to both p and q, contrary to the hypothesis that $p \ncong q$. $\square$

Corollary 20 *If p and q are non-isomorphic, selective ultrafilters on $\mathbb{N}$ and if U is partitioned into finitely many pieces, then there are sets $A \in p$ and $B \in q$ such that $(A \times B) \cap U$ is included in one of the pieces.*

Proof Partition $\mathbb{N}^2$ into

- the finitely many given pieces of U,
- their images under τ, giving a partition of D, and
- Δ.

By the preceding corollary and Lemma 5, find $A \in p$ and $B \in q$ such that $A \times B$ meets only two of these pieces. Since it must meet both U and D (as A and B are infinite), it meets only one of the pieces of U. $\square$

Remark 21 A theorem of Kunen, first published in [3], says that selective ultrafilters on $\mathbb{N}$ are the same as *Ramsey* ultrafilters, i.e., those $p \in \beta\mathbb{N}$ such that, whenever $[\mathbb{N}]^2$ is partitioned into finitely many pieces, there is $A \in p$ such that $[A]^2$ lies in one of the pieces. (In other words, the homogeneous set in Ramsey's theorem for pairs [14, 7] can be taken to lie in p.) By Lemma 9, this is the same as $|\widetilde{\iota}^{-1}[\{(p,p)\}]| = 3$. One direction of Kunen's theorem is very easy; given $f : \mathbb{N} \to \mathbb{N}$, partition $[N]^2$ into $\{\{x,y\} : f(x) = f(y)\}$ and its complement, and notice that, if $[A]^2$ is included in one piece, then f is constant or one-to-one on A. The (non-trivial) converse direction can be proved by a minor modification of the last half of the proof of the preceding theorem. We sketch this argument.

Let $p \in \mathbb{N}^*$ be selective. We intend to show that the only ultrafilter $s \in \widetilde{\iota}^{-1}[\{(p,p)\}]$ that contains U is $p \otimes p$. Proceed exactly as in the proof of Theorem 17 starting at the place "Temporarily suppose $U \in s$", replacing q by p everywhere. Continue to the point where we find that $\widetilde{g_1 \circ f}(p) = \widetilde{g_2 \circ g}(p)$. As before, the right side of this equation is non-principal, since $g_2 \circ g$ is finite-to-one on a set in p, but this is no longer a contradiction, since we no longer have the hypothesis of "no common minorant in $\leq_{RK}$." Instead, we argue as follows. Since p is selective, $g_2 \circ g$ must be one-to-one on a set in p, and it agrees, on a set in p, with a bijection $e : \mathbb{N} \to \mathbb{N}$ (as in the proof of Lemma 3). Then from $\widetilde{g_1 \circ f}(p) = \widetilde{g_2 \circ g}(p) = \widetilde{e}(p)$ we infer that $\widetilde{e^{-1} \circ g_1 \circ f}(p) = p$. By Lemma 2, this means that $e^{-1} \circ g_1 \circ f$ is the identity map on some set in p, and therefore that $g_1 \circ f = e$ on a set in p.

But we also had, in the proof of Theorem 17, that $g_1 \circ f \circ \pi_1 = g_2 \circ g \circ \pi_2$ on a set in s. Since both $g_1 \circ f$ and $g_2 \circ g$ agree with e on sets in $p = \widetilde{\pi}_1(s) = \widetilde{\pi}_2(s)$, it follows that $\{(x,y) : e(x) = e(y)\} \in s$. As e is a bijection, this means $\Delta \in s$, contrary to our assumption that $U \in s$. This contradiction completes the proof of Kunen's theorem.

We remark that Kunen's original proof is more direct. The purpose of our proof here is not to try to improve on his proof but to show the connection between his result and the methods used in the present paper.

The following result comes from combining the preceding remark and Corollaries 4 and 19

Corollary 22 *Let p and q be selective ultrafilters on $\mathbb{N}$. Then $|\widetilde{\iota}^{-1}[\{(p,q)\}]|$ is 3 if they are isomorphic and 2 otherwise.*

Theorem 23 *If $p, q \in \mathbb{N}^*$ and $|\widetilde{\iota}^{-1}[\{(p,q)\}]| \leq 5$, then at least one of p, q is a P-point.*

Proof Suppose neither p nor q is a P-point. By Lemma 11, let $f, g : \mathbb{N} \to \mathbb{N}$ be functions such that f is neither finite-to-one nor constant on any set in p, g is neither finite-to-one nor constant on any set in q, and $f(x), g(x) \leq x$ for all x. By Lemma 5, it suffices to show that, for each $A \in p$ and each $B \in q$, the product $A \times B$ meets all six of the sets

$$\begin{aligned}
C_1 &= \{(x,y) : y > x > g(y) > f(x)\} \\
C_2 &= \{(x,y) : y > x \geq f(x) > g(y)\} \\
C_3 &= \{(x,y) : y \geq g(y) > x \geq f(x)\} \\
C_4 &= \{(x,y) : x > y > f(x) > g(y)\} \\
C_5 &= \{(x,y) : x > y \geq g(y) > f(x)\} \\
C_6 &= \{(x,y) : x \geq f(x) > y \geq g(y)\}.
\end{aligned}$$

By symmetry, it suffices to show that $A \times B$ meets each of C_1, C_2, and C_3.

For C_1, first recall that f is not finite-to-one on any set in p, in particular on A. So pick n such that $A' = \{x \in A : f(x) = n\}$ is infinite. Since g is not constant on any set in q, it cannot be bounded above on any set in q, in particular on B, so $B' = \{y \in B : g(y) > n\} \in q$. Since g is not finite-to-one on B', pick m such that $B'' = \{y \in B' : g(y) = m\}$ is infinite. Then, since A' is infinite, pick $x \in A'$ with $x > m$, and since B'' is infinite, pick $y \in B''$ with $y > x$. Then we have $y > x > m = g(y) > n = f(x)$, and so $(x,y) \in (A \times B) \cap C_1$.

For C_2, pick n such that $B' = \{y \in B : g(y) = n\}$ is infinite. Then, since f is not bounded above on any set from p, pick $x \in A$ with $f(x) > n$. Finally, pick $y \in B'$ with $y > x$. We have $y > x \geq f(x) > n = g(y)$, and so $(x,y) \in (A \times B) \cap C_2$.

Finally, for C_3 begin by picking any $x \in A$. Since g isn't bounded above on any set from q, in particular it isn't bounded above by x on B. So pick $y \in B$ with $g(y) > x$. Then we have $y \geq g(y) > x \geq f(x)$, and so $(x,y) \in (A \times B) \cap C_3$. □

Theorem 24 *If $p \in \mathbb{N}^*$ and $|\widetilde{\iota}^{-1}[\{(p,p)\}]| \leq 8$ then p is a P-point.*

Observe that, in view of Lemma 8, the hypothesis of this theorem is equivalent to $|\widetilde{\iota}^{-1}[\{(p,p)\}]| \leq 7$. By Lemma 9, it is also equivalent to requiring that, whenever $[\mathbb{N}]^2$ is partitioned into finitely many pieces, then there is

$A \in p$ such that $[A]^2$ meets at most 3 of the pieces. In this form, the theorem was stated and proved model-theoretically in [2, Theorem4].

Proof Suppose p is not a P-point. By Lemma 11, fix a function $f : \mathbb{N} \to \mathbb{N}$ such that $f(x) \leq x$ for all x and f is neither finite-to-one nor constant on any set from p. By Lemma 9, it suffices to show that, for every $A \in p$, the square $A \times A$ meets all four of the following disjoint subsets of U:

$$\begin{aligned} C_1 &= \{(x,y) \in \mathbb{N}^2 : y > x \geq f(x) = f(y)\} \\ C_2 &= \{(x,y) \in \mathbb{N}^2 : y > x \geq f(x) > f(y)\} \\ C_3 &= \{(x,y) \in \mathbb{N}^2 : y > x > f(y) > f(x)\} \\ C_4 &= \{(x,y) \in \mathbb{N}^2 : y \geq f(y) > x \geq f(x)\}. \end{aligned}$$

For C_1, pick n such that $A' = \{x \in A : f(x) = n\}$ is infinite. Then pick $x < y$ both in A' and observe that $y > x \geq f(x) = n = f(y)$ and so $(x,y) \in A^2 \cap C_1$.

For C_2, again pick n such that $A' = \{x \in A : f(x) = n\}$ is infinite. As f is not constant on any member of p, it cannot be bounded above on any member of p; in particular, $A'' = \{x \in A : f(x) > n\} \in p$. Pick any $x \in A''$ and then pick any $y > x$ in A'. We have $y > x \geq f(x) > n = f(y)$ and so $(x,y) \in A^2 \cap C_2$.

For C_3, again pick n such that $A' = \{x \in A : f(x) = n\}$ is infinite and again notice that $A'' = \{x \in A : f(x) > n\} \in p$. Pick m such that $A''' = \{x \in A'' : f(x) = m\}$ is infinite. Notice that, by definition of A'', $m > n$. Pick $x > m$ in A'; then pick $y > x$ in A'''. Then we have $y > x > m = f(y) > n = f(x)$ and so $(x,y) \in A^2 \cap C_3$.

Finally, for C_4, pick any $x \in A$. As f is not bounded above on any set in p, we have that $A' = \{y \in A : f(y) > x\}$ is in p and is therefore infinite. Pick any $y \in A'$. Then we have $y \geq f(y) > x \geq f(x)$ and so $(x,y) \in A^2 \cap C_4$. □

5 Tensor Products and Optimality

In this section, we show that the cardinality bounds of 2, 5, and 8 in Theorems 17, 23, and 24 are optimal provided enough selective ultrafilters exist. The last of these was already done in [2] with an elementary proof (modulo the use of Corollary 20 above), so we do not repeat the proof here but we state the result for completeness.

Theorem 25 *Assume that there are infinitely many, pairwise non-isomorphic, selective ultrafilters on* $\mathbb{N}$. *Then there is a non-P-point* $p \in \mathbb{N}^*$ *such that* $|\widetilde{\iota}^{-1}[\{(p,p)\}]| = 9$.

Proof See [2, Example 2]. $\square$

For the rest of the results of this section, we shall need some additional information about tensor products of ultrafilters. Since we shall need to construct non-P-points, the following lemma will be a useful tool. The non-P-points we need will be obtained as isomorphic copies on $\mathbb{N}$ of certain tensor product ultrafilters on $\mathbb{N}^2$.

Lemma 26 *If* $p, q \in \mathbb{N}^*$ *then* $p \otimes q$ *is not isomorphic to a P-point.*

Proof The projection π_1 is neither constant nor finite-to-one on any set in $p \otimes q$. Indeed, if π_1 is constant on a set $A \subseteq \mathbb{N}^2$ then all but one of the slices A_x will be empty, while if π_1 is finite-to-one on A then all the slices are finite. In either case, it cannot be true that almost all (with respect to p) of the slices are in q, since both p and q are non-principal. $\square$

Lemma 27 *Suppose* $s \leq_{RK} p \otimes q$ *where* $p, q, s \in \mathbb{N}^*$ *and* s *is a P-point. Then* $s \leq_{RK} p$ *or* $s \leq_{RK} q$.

Proof Fix a function $f : \mathbb{N}^2 \to \mathbb{N}$ such that $s = \widetilde{f}(p \otimes q)$. For each $n \in \mathbb{N}$, define $f_n : \mathbb{N} \to \mathbb{N}$ by $f_n(y) = f(n, y)$, and consider the ultrafilters $\widetilde{f_n}(q)$.

We claim that, for every set $A \in p$, the closure in $\beta\mathbb{N}$ of $\{\widetilde{f_n}(q) : n \in A\}$ contains s. To see this, consider any basic open neighborhood of s; it has the form $\overline{B}$ for some $B \in s$. Since $s = \widetilde{f}(p \otimes q)$, the ultrafilter $p \otimes q$ must contain $f^{-1}[B]$. That means that, for almost all n with respect to p, the slice $(f^{-1}[B])_n$ belongs to q. But that slice is exactly ${f_n}^{-1}[B]$. So we have an $n \in A$ (in fact many of them) such that ${f_n}^{-1}[B] \in q$ and thus $\widetilde{f_n}(q) \in \overline{B}$. This establishes our claim.

On the other hand, we claim also that s, being a P-point, cannot be in the closure of any countable subset $\{r_n : n \in \mathbb{N}\}$ of $\mathbb{N}^* - \{s\}$. Indeed, given such a subset, we could find, for each $n \in \mathbb{N}$, a set $C_n \in s$ such that $C_n \notin r_n$. As s is a P-point, we could find $C \in s$ such that $C - C_n$ is finite for each n. As r_n is non-principal, we would have $C \notin r_n$. Thus, $\overline{C}$ would be an open neighborhood of s containing no r_n. This establishes the second claim.

Comparing the two claims, we find that $\{n : \widetilde{f_n}(q) \in \mathbb{N}^* - \{s\}\}$ cannot be in p. Being an ultrafilter, p must therefore contain either $\{n : \widetilde{f_n}(q) \text{ principal}\}$ or $\{n : \widetilde{f_n}(q) = s\}$.

In the first case, we have a set $D \in p$ and, for each $n \in D$, a set $E_n \in q$ on which f_n is constant, say with value $g(n)$. Then f coincides with $g \circ \pi_1$ on the set $\{(x, y) : x \in D, y \in E_x\} \in p \otimes q$. So

$$s = \widetilde{f}(p \otimes q) = \widetilde{g}(\widetilde{\pi}_1(p \otimes q)) = \widetilde{g}(p) \leq_{RK} p.$$

In the second case, we have an n (in fact many of them) for which

$$s = \widetilde{f_n}(q) \leq_{RK} q.$$

$\square$

The main tool used in our constructions in this section is a description of all the ultrafilters on $\mathbb{N}^n$ that project to n given, pairwise non-isomorphic, selective ultrafilters on $\mathbb{N}$. The first and largest step toward this description is the following theorem.

Theorem 28 *Suppose s is an ultrafilter on $\mathbb{N}^n$ containing the set*

$$U_n = \{(x_1, \ldots, x_n) \in \mathbb{N}^n : x_1 < x_2 < \cdots < x_n\}.$$

Suppose further that its images under the projection maps, $p_i = \widetilde{\pi}_i(s)$, are n pairwise non-isomorphic, selective ultrafilters. Then $s = p_1 \otimes p_2 \otimes \cdots \otimes p_n$.

Proof We proceed by induction on n. The case $n = 1$ is trivial if we use the obvious interpretation of the "tensor product of one ultrafilter" as that ultrafilter. This interpretation agrees with the equation

$$p_1 \otimes \cdots \otimes p_n = (p_1 \otimes \cdots \otimes p_{n-1}) \otimes p_n,$$

which we shall use below. So it is a legitimate basis for our induction.

For the induction step, assume that we are given s and p_i's as in the hypothesis of the theorem, and assume that the theorem has been proved with $n-1$ in place of n. We write $\lambda : \mathbb{N}^n \to \mathbb{N}^{n-1}$ for the projection from $\mathbb{N}^n$ to the product of the first $n-1$ factors. Thus $\pi_i \circ \lambda = \pi_i$ for $1 \leq i < n$.

Notice that $\widetilde{\lambda}(s)$ is an ultrafilter on $\mathbb{N}^{n-1}$, that it contains U_{n-1}, and that its projections are the p_i for $1 \leq i < n$. So the induction hypothesis tells us that

$$\widetilde{\lambda}(s) = p_1 \otimes \cdots \otimes p_{n-1}.$$

What we must prove, therefore, is $s = \widetilde{\lambda}(s) \otimes p_n$. Since $p_n = \pi_n(s)$ and since λ and π_n are the two projection maps of the product $\mathbb{N}^{n-1} \times \mathbb{N}$ (i.e., $\mathbb{N}^n$ viewed as a product of two factors), Theorem 13 reduces our task to verifying the following claim:

If $f : \mathbb{N}^{n-1} \to \mathbb{N}$ and $g : \mathbb{N} \to \mathbb{N}$ then either g is constant on a set in p_n or

$$\{(x_1, \dots, x_{n-1}, x_n) : f(x_1, \dots, x_{n-1}) < g(x_n)\} \in s.$$

To prove this claim, let f and g be given, and assume that g is not constant on any set in p_n. As p_n is selective, g is one-to-one on some set $A \in p_n$. We define a function $h : \mathbb{N} \to \mathbb{N}$ as follows. If $k \notin g[A]$ then define $h(k)$ arbitrarily. If $k \in g[A]$, then let a be the unique element of A with $g(a) = k$ (it's unique because g is one-to-one on A) and consider all $(n-1)$-tuples $(x_1, \dots, x_{n-1})$ whose components are all $\leq a$. As there are only finitely many such tuples, let $h(k)$ be a number bigger than all the corresponding values of $f(x_1, \dots, x_{n-1})$. Thus, for any n-tuple $(x_1, \dots, x_n) \in U_n \cap {\pi_n}^{-1}[A]$, we have $f(x_1, \dots, x_{n-1}) < h(g(x_n))$.

Now suppose, toward a contradiction, that the conclusion of the claim failed, so, since s is an ultrafilter, we would have $g(x_n) \leq f(x_1, \dots, x_{n-1})$ on a set in s. But we have just shown that $f(x_1, \dots, x_{n-1}) < h(g(x_n))$ on a set in s, namely $U_n \cap {\pi_n}^{-1}[A]$. So Lemma 18 applies, with $\mathbb{N}^{n-1}$ in the role of X, with $\mathbb{N}$ in the role of Y, with s in the role of r, with $f \circ \lambda$ and $g \circ \pi_n$ in the roles of f_1 and f_2, and with h and $\mathrm{id}_{\mathbb{N}}$ in the roles of h_1 and h_2. That lemma provides finite-to-one functions $g_1, g_2 : \mathbb{N} \to \mathbb{N}$ such that $g_1 \circ f \circ \lambda = g_2 \circ g \circ \pi_n$ on a set in s. In particular,

$$\widetilde{g_2 \circ g}(p_n) = \widetilde{g_2 \circ g \circ \pi_n}(s) = \widetilde{g_1 \circ f \circ \lambda}(s) = \widetilde{g_1 \circ f}(p_1 \otimes \cdots \otimes p_{n-1}).$$

The ultrafilter r described by the left side of this equation is isomorphic to p_n, since it isn't principal (remember g is one-to-one on a set in p_n and g_2 is finite-to-one) and p_n is selective. In particular r is a P-point. Now looking at the right side of the equation and repeatedly applying Lemma 27, we find that r is also $\leq_{RK} p_i$ for some i in the range $1 \leq i < n$. As p_i is selective and r isn't principal, we infer that $r \cong p_i$. So we have r isomorphic to both p_n and another p_i. This is absurd, as $p_i \ncong p_n$ by hypothesis. This contradiction completes the proof of the claim and thus of the theorem. $\square$

We shall need a description of all the ultrafilters $s \in \beta(\mathbb{N}^n)$ satisfying $\widetilde{\pi}_i(s) = p_i$ for given, pairwise non-isomorphic, selective ultrafilters $p_1, \dots, p_n$.

The preceding theorem provides such a description when s contains U_n. The general case will be an easy consequence once we set up appropriate notation.

Definition 29 Let σ be any permutation of $\{1,\ldots,n\}$. The induced bijection $\hat{\sigma}:\mathbb{N}^n\to\mathbb{N}^n$ is defined by

$$\hat{\sigma}(x_1,\ldots,x_n)=(x_{\sigma^{-1}(1)},x_{\sigma^{-1}(2)},\ldots,x_{\sigma^{-1}(n)}).$$

The induced self-homeomorphism of $\beta(\mathbb{N}^n)$ will be denoted by $\widetilde{\sigma}$; that is we use only a tilde where strictly speaking we should have both a hat and a tilde.

If the use of σ^{-1} instead of just σ in the definition seems an unnecesary complication, the reader should check that (1) our definition does what would intuitively be described as "permuting the components of an n-tuple according to σ" and (2) it satisfies $\widehat{\sigma_1\circ\sigma_2}=\hat{\sigma}_1\circ\hat{\sigma}_2$ whereas a definition omitting the inversions would reverse the order of composition.

The definition can be usefully rewritten as $\pi_i\circ\hat{\sigma}=\pi_{\sigma^{-1}(i)}$.

Notice that $\hat{\sigma}[U_n]=\{(x_1,\ldots,x_n):x_{\sigma(1)}<\cdots<x_{\sigma(n)}\}$.

Corollary 30 *Suppose s is an ultrafilter on $\mathbb{N}^n$ whose projections $\widetilde{\pi}_i(s)=p_i$ are n pairwise non-isomorphic, selective ultrafilters. Then there is a unique permutation σ of $\{1,\ldots,n\}$ such that $\hat{\sigma}[U_n]\in s$. For this σ we have*

$$s=\widetilde{\sigma}(p_{\sigma(1)}\otimes\cdots\otimes p_{\sigma(n)}).$$

Proof The sets $\hat{\sigma}[U_n]$ for different σ's are disjoint, so s contains at most one of them. Furthermore, these sets and the sets $E_{ij}=\{(x_1,\ldots,x_n):x_i=x_j\}$ for $1\le i<j\le n$ cover $\mathbb{N}^n$, so s must contain one of them. Since $\pi_i=\pi_j$ on E_{ij} and since $\widetilde{\pi}_i(s)\neq\widetilde{\pi}_j(s)$, s cannot contain any E_{ij} so it must contain one of the $\hat{\sigma}[U_n]$'s. So we have $\hat{\sigma}[U_n]\in s$ for exactly one σ, which will be fixed for the rest of the proof.

From $\hat{\sigma}[U_n]\in s$ and the observation that "hat commutes with inverse," we obtain that $U_n\in\widetilde{\sigma^{-1}}(s)$. Furthermore, the projections of $\widetilde{\sigma^{-1}}(s)$ are

$$\widetilde{\pi}_i(\widetilde{\sigma^{-1}}(s))=\widetilde{\pi_{\sigma(i)}}(s)=p_{\sigma(i)}.$$

So Theorem 28 applies to $\widetilde{\sigma^{-1}}(s)$ and tells us that

$$\widetilde{\sigma^{-1}}(s)=p_{\sigma(1)}\otimes\cdots\otimes p_{\sigma(n)}.$$

Applying $\widetilde{\sigma}$ to both sides, we get the conclusion of the corollary. $\square$

Armed with these tools, we can now establish the optimality of the numbers in Theorems 17 and 23.

Theorem 31 *Assume that there are at least three non-isomorphic, selective ultrafilters on* $\mathbb{N}$. *Then there are ultrafilters* $p, q \in \mathbb{N}^*$ *such that* $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 3$ *and* p *is not a P-point.*

Of course the q here must be a P-point, by Theorem 23.

Proof Let p_1, p_2, and p_3 be non-isomorphic, selective ultrafilters. We take p to be an isomorphic copy on $\mathbb{N}$ of the ultrafilter $p_1 \otimes p_2$ on $\mathbb{N}^2$. We take q to be p_3. By Lemma 26, p is not a P-point. To establish that $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 3$, it suffices to show that there are only three ultrafilters s on $\mathbb{N}^3$ satisfying $\widetilde{\lambda}(s) = p_1 \otimes p_2$ and $\widetilde{\pi}_3(s) = p_3$, where, as before, $\lambda : \mathbb{N}^3 \to \mathbb{N}^2$ is the projection to the first two components. (In detail: If $f : \mathbb{N}^2 \to \mathbb{N}$ is a bijection sending $p_1 \otimes p_2$ to p, then $f'(x,y,z) = (f(x,y),z)$ is a bijection $\mathbb{N}^3 \to \mathbb{N}^2$ sending the ultrafilters s considered here to the ultrafilters in $\widetilde{\iota}^{-1}[\{(p,q)\}]$.)

Any such s has

$$\widetilde{\pi}_1(s) = \widetilde{\pi_1 \circ \lambda}(s) = \widetilde{\pi}_1(p_1 \otimes p_2) = p_1$$

and similarly $\widetilde{\pi}_2(s) = p_2$. So s is as described in Corollary 30 for some σ. Of the six permutations σ of $\{1,2,3\}$, only three give an s satisfying $\widetilde{\lambda}(s) = p_1 \otimes p_2$. (The other three give $\widetilde{\tau}(p_2 \otimes p_1)$ instead.) In other words, since $p_1 \otimes p_2$ contains $U = \{(x,y) : x < y\}$, s must contain $\lambda^{-1}[U]$, which meets $\hat{\sigma}[U_3]$ for only three of the six permutations σ, namely those for which $\sigma^{-1}(1) < \sigma^{-1}(2)$. So s must be the ultrafilter described in Corollary 30 for one of these three σ's. $\square$

Theorem 32 *Assume that there exist four pairwise non-isomorphic, selective ultrafilters on* $\mathbb{N}$. *Then there are two non-P-points* $p, q \in \mathbb{N}^*$ *such that* $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 6$.

Proof Let p_1, p_2, p_3 and p_4 be pairwise non-isomorphic, selective ultrafilters. Let p and q be isomorphic copies on $\mathbb{N}$ of $p_1 \otimes p_2$ and $p_3 \otimes p_4$, respectively. So by Lemma 26, neither of them is a P-point. As in the preceding proof, to show that $|\widetilde{\iota}^{-1}[\{(p,q)\}]| = 6$, it suffices to show that there are only 6 ultrafilters s on $\mathbb{N}^4$ satisfying $\widetilde{\mu}(s) = p_1 \otimes p_2$ and $\widetilde{\nu}(s) = p_3 \otimes p_4$, where μ and

ν are the projection functions $\mathbb{N}^4 \to \mathbb{N}^2$ defined by $\mu(x_1, x_2, x_3, x_4) = (x_1, x_2)$ and $\nu(x_1, x_2, x_3, x_4) = (x_3, x_4)$. Also as in the preceding proof, any such s must satisfy

$$\widetilde{\pi}_1(s) = \widetilde{\pi_1 \circ \mu}(s) = \widetilde{\pi}_1(p_1 \otimes p_2) = p_1$$

and similarly $\widetilde{\pi}_i(s) = p_i$ for the remaining three values of i. Thus, Corollary 30 applies and describes s completely in terms of a permutation σ of $\{1, 2, 3, 4\}$. There are 24 such permutations, but only half produce an s with $\widetilde{\mu}(s) = p_1 \otimes p_2$, and only half of these produce an s with $\widetilde{\nu}(s) = p_3 \otimes p_4$. Specifically, the condition on $\widetilde{\mu}(s)$ requires $\sigma^{-1}(1) < \sigma^{-1}(2)$ and the condition on $\widetilde{\nu}(s)$ requires $\sigma^{-1}(3) < \sigma^{-1}(4)$. These requirements are satisfied by exactly 6 of the 24 permutations, so there are 6 possible ultrafilters s. $\square$

Remark 33 The proofs of the last two theorems admit the following common generalization. Suppose $p_1, \ldots, p_n$ are n pairwise non-isomorphic, selective ultrafilters. Let m be a number in the range $1 \leq m < n$, let p be an isomorphic copy on $\mathbb{N}$ of $p_1 \otimes \cdots \otimes p_m$, and let q be an isomorphic copy on $\mathbb{N}$ of $p_{m+1} \otimes \cdots \otimes p_n$. Then the cardinality of $\widetilde{\iota}^{-1}[\{(p, q)\}]$ is the binomial coefficient $\binom{n}{m}$.

Remark 34 The techniques presented in this paper can be straightforwardly extended to apply to the natural maps $\theta : \beta(\mathbb{N}^k) \to (\beta\mathbb{N})^k$ for all finite $k \geq 2$. For example, one can show that, if

$$|\theta^{-1}[\{(p_1, \ldots, p_k)\}]| < k! \cdot 1 \cdot 3 \cdot \cdots \cdot (2k - 1),$$

then at least one of the p_i is a P-point. If there are at least $2k$ non-isomorphic selective ultrafilters, then by taking each p_i to be the product of two selective ultrafilters, all factors of all the p_i's being non-isomorphic, we obtain non-P-points satisfying

$$|\theta^{-1}[\{(p_1, \ldots, p_k)\}]| = k! \cdot 1 \cdot 3 \cdot \cdots \cdot (2k - 1).$$

References

[1] A. Blass, *Orderings of Ultrafilters*, Ph. D. thesis, Harvard University (1970).

[2] A. Blass, *Amalgamation of nonstandard models of arithmetic*, J. Symbolic Logic 42 (1977) 372–385.

[3] D. Booth, *Ultrafilters on a countable set*, Ann. Math. Logic 2 (1970) 1–24.

[4] P. E. Cohen, *P-points in random universes*, Proc. Amer. Math. Soc. 74 (1979) 318–321.

[5] W. W. Comfort and S. Negrepontis, *The Theory of Ultrafilters*, Springer-Verlag (1974).

[6] M. Daguenet, *Rapport entre l'ensemble des ultrafiltres admettant un ultrafiltre donné pour image et l'ensemble des images de cet ultrafiltre*, Comm. Math. Univ. Carolinae 16 (1975) 99–113.

[7] R. Graham, B. Rothschild, and J. Spencer, *Ramsey Theory*, Wiley (1977).

[8] N. Hindman, *On P-like spaces and their product with P-spaces*, Ph. D. thesis, Wesleyan University (1969).

[9] N. Hindman, *Preimages of points under the natural map from $\beta(\mathbb{N}\times\mathbb{N})$ to $\beta\mathbb{N}\times\beta\mathbb{N}$*, Proc. Amer. Math. Soc. 37 (1973) 603–608.

[10] N. Hindman and D. Strauss, *Algebra in the Stone-Čech Compactification*, W. de Gruyter (1998).

[11] K. Kunen, *Some points in βN*, Math. Proc. Cambridge Phil. Soc. 80 (1976) 385–398.

[12] G. Moche, *Preimages of points under the natural map from $K(\beta(\mathbb{N}\times\mathbb{N}))$ to $K(\beta\mathbb{N})\times K(\beta\mathbb{N})$*, Ph.D. thesis, Howard University (2002).

[13] C. Puritz, *Skies, constellations and monads*, in Contributions to Non-Standard Analysis, W. A. J. Luxemburg and A. Robinson, eds., Studies in Logic and Foundations of Math. 69 (1972), North-Holland, pp. 215–243.

[14] F. P. Ramsey, *On a problem of formal logic*, Proc. London Math. Soc. 30 (1930) 264–286.

[15] S. Shelah, *Proper Forcing*, Springer-Verlag, Lecture Notes in Mathematics 940 (1982)