Methods for strong human germline engineering

Tsvi Benson-Tilsen

3 March 2025

Contents

Introduction **3**

Takeaways **6**

Disclaimers **7**

Terms **8**
- Non-standard terms . . . 8
- Standard terms . . . 8

Reproductive genomic vectoring **8**
- Genomic vectoring (GV) . . . 9
- Strong GV and why it matters . . . 9
- Polygenic scores . . . 10
- Editing and selection . . . 10
- Other types of genomic vectoring (GV) . . . 11

Comparing editing and selection **11**
- Editing produces out-of-distribution DNA . . . 11
- Editing is already minimally working technology . . . 11
- Headwind of mutation . . . 11
- Fixing de novo mutations . . . 12
- Bypassing natural integrity checks . . . 12
- Selection makes similar children . . . 12
- Selection and editing benefit differently from stronger PGSes . . . 13
- Selection distributes Euclideanly, editing distributes Hammingly . . . 14
- Selection beats editing, then editing beats selection . . . 14
- Editing can reach more places than selection . . . 15
- The ceiling of safe vectoring . . . 15
- Many GV methods synergize . . . 15

Reproductive GV and epigenomic correctness (EC) **16**
- The epigenomic correctness problem . . . 16
- The near-miss hazard of epigenomic correction . . . 18
- Going from 0.01 to 0.99 . . . 21
- Biology's *scopes problem . . . 22

Methods to handle epigenomic correctness **23**
- In vitro gametogenesis (IVG) . . . 24
  - The basic elements of gametogenesis . . . 25
  - IVG is a higher bar than epigenomic correction . . . 27
  - Other general remarks about IVG . . . 28
  - The general state of IVG research . . . 29
  - In vitro oogenesis (IVO) . . . 37
  - In vitro spermatogenesis (IVS) . . . 49
- Using natural epigenomic correction . . . 57
  - No clear boundary between natural and artificial EC-making . . . 58

The timing problem with natural EC-making . . . 59
Alternative gonadal tissue . . . 59
Using natural reproductive DNA . . . 61
Natural EC interrupt methods . . . 61
Imprint maintenance problems . . . 61
Epigenetic CRISPR editing . . . 62
Donor embryo . . . 64
Hybrid natural and artificial EC-making . . . 64

How GV and EC interact **65**
A simple-ish and wrong model of combining GV and EC . . . 65
GV and EC are not conceptually separable . . . 65
Features of GV and EC methods that affect compatibility . . . 66
How feasible are different epigenomic correctness methods? . . . 67
How epigenomically disruptive are different GV methods? . . . 68
No epigenomic disruption . . . 68
Little epigenomic disruption . . . 68
Significant epigenomic disruption . . . 68
Haploid vs. diploid . . . 69

Summary of genomic vectoring methods **69**

Method: Simple embryo selection **72**
The power of simple embryo selection . . . 72
Super-duper-ovulation . . . 75

Method: Gamete selection **76**
Single gamete selection . . . 76
Gamete selection is mostly hypothetical . . . 76
But sperm selection might be doable . . . 76
The power of single gamete selection . . . 76
Double gamete selection . . . 77
The power of double gamete selection . . . 77
The limits of extreme gamete selection . . . 78

Method: Chromosome selection **79**
Implementation of chromosome selection . . . 80
Mechanical manipulation . . . 80
Whole cell fusion . . . 81
The power of simple chromosome selection . . . 81
Selecting a haploid genome . . . 81
Selecting a diploid genome . . . 82
Selecting an epigenomically correct haploid . . . 83
Selecting average chromosomes . . . 84
Selecting a diploid source of haploids . . . 85
The power of recombinant chromosome selection . . . 85
One-donor recombinant chromosome selection . . . 86
Many-donor recombinant chromosome selection . . . 86
Double recombinant chromosome selection . . . 86
Fractional haploid donation . . . 87

Method: Iterated recombinant selection **89**
Iterated embryo selection . . . 90
Iterated meiotic selection . . . 90

Auxiliary: Enhancements to meiotic methods **91**

Method: Iterated multiplex CRISPR editing **93**
DNA damage from editing . . . 94
Naive ESC editing . . . 94

Hulk sperm . . . 94
Magic rainbow sperm . . . 95

Conclusion **95**

Appendix: In vitro spermatogenesis studies **95**

Appendix: Cheap DNA segment sensing **97**
The problem of cheap DNA sensing . . . 97
Basic chromosome sensing . . . 97
Sensing crossovers . . . 97
An obstacle with sperm . . . 98
Cheap sensing at scale with FACS . . . 98
Sensing chromosome index . . . 98

Appendix: Best crossover **99**
What does recombination do to DNA? . . . 99
Math analysis of crossovers . . . 102
Simulation results of crossovers . . . 104
Refinements to the crossover model . . . 106
Non-uniform crossover points . . . 106
Multiple crossovers . . . 106

Appendix: The costs of iterated meiotic selection **108**
The power of segmented selection . . . 108
Breakdown at high granularity . . . 109
Segmented selection with one or more donors . . . 110
The cost of poor man's chromosome selection . . . 110
The cost of single-chromosome IMS . . . 111

Appendix: Detailed estimation of embryo selection **114**
Simple embryo selection simulations . . . 114
Analysis of simple embryo selection . . . 114
What is the distribution of embryos from one couple? . . . 114
What is the maximum scoring embryo from one couple? . . . 117

Appendix: Variation in chromosome length **123**
The issue with chromosome lengths . . . 123
Do chromosome lengths matter much? . . . 124
Why don't chromosome lengths matter much? . . . 126
Ignoring the 23rd chromosome . . . 127

Acknowledgements **128**

License **128**

Introduction

This article summarizes the technical pathways to make healthy humans with significantly modified genomes. These are the pathways that I'm aware of and that seem plausibly feasible in the next two decades. A short summary, in a diagram:

	know what your target looks like	make a cell with the target
genome	**polygenic scores (PGS)** constructed by genome-wide association studies and other analysis. requires phenotype+genotype data. clinicians inform parents of the meaning of available PGSes.	**genomic vectoring (GV)** editing or selection. selection methods: embryo selection, gamete selection, chromosome selection, iterated meiotic selection.
epigenome	**epigenomic correctness** epigenomic atlases of natural reproductive DNA (gametes, embryos). bisulfite seq, scRNA seq, scATAC seq, nanopore, ChIP-seq.	**epigenomic correction** in vitro EC-making; using natural reproductive DNA; in vivo reprogramming.

Berkeley Genomics Project
2025 license: CC BY-SA 4.0

Annotated table of contents:

- Reproductive genomic vectoring explains the general idea of human germline genomic engineering, and distinguishes editing and selection.
- Comparing editing and selection talks about general differences between the two kinds of genomic vectoring methods.
- Reproductive GV and epigenomic correctness (EC), Methods to handle epigenomic correctness, and How GV and EC interact discuss the epigenomic correctness problem in germline engineering—what it is, why it matters, and how to address it.
- Summary of genomic vectoring methods gives an annotated table of contents for the following Methods sections. The Methods sections—on Simple embryo selection, Gamete selection, Chromosome selection, Iterated recombinant selection, and Iterated multiplex CRISPR editing—give more detail about each genomic vectoring method: what it is, obstacles, variations, and how powerful it is.
- The appendices give additional technical information, if you're looking around and saying "I'm not in the weeds *enough*, I want to be *more* in the weeds.".

Here's a sneak peek about the strength of different genomic vectoring methods:

Strength comparison of methods for strong human germline engineering

March 2025 | Berkeley Genomics Project | license: CC BY-SA 4.0 https://creativecommons.org/licenses/by-sa/4.0/

k	M(k)	M(2k)
1	0	0.56
2	0.56	1.03
4	1.03	1.42
10	1.53	1.87
23	1.93	2.22
300	2.87	3.09
20000	4.01	4.17

```
import numpy as np
def expected_max(test_runs=1_000, k=2):
  X = np.random.randn(test_runs, k)
  return np.mean(np.max(X, axis=1))"
```

M(k) means "the expected value, in SDs, of the largest of k samples from a standard gaussian"

10 embryo selection gives 1.53/√2=1 σ, 6 IQ points; haploid is .7 σ or 4 IQ

kman's editing estimates: https://www.lesswrong.com/posts/DfrSZaf3JC8vJdbZL/how-to-make-superbabies

for editing rows, column shows SDs for diploid editing, from kman's estimates based on genomics data that corresponds to a .35 PGS for IQ

methods organized in order of my very non-rigorous guess at how fast they *could* be doable if seriously tried

assuming R=.4 PGS

method name	prerequisites	formula for raw SDs	parameters	expected raw SDs	expected IQ points	notes
simple embryo selection	already doable	M(embryos)/√2	5 embryos	0.8 ± 0.6	5 ± 4	
			10 embryos	1.1 ± 0.5	6 ± 3	
			100 embryos	1.7 ± 0.4	10 ± 2	
			1000 embryos	2.2 ± 0.4	14 ± 2	
ESC CRISPR editing, small count	weak ESC imprinting maintenance (for a week, say); ESC SCNT works	6 + SDs × 15 from kman. diploid.	~15 edits?		14	0.56
sperm average-chromosome selection (+ simple embryo selection, 10 embryos)	sensing chromosome index chromosome transplantation sperm DNA must be mechanically robust enough	.7 + [M(donors)/√2] × (22/√46)	1 donor	0.7	4	
			2 donors	2	12	
			4 donors	3.1	19	
			10 donors	4.2	25	
			23 donors	5.1	31	
			300 donors	7.3	44	
			20,000 donors	9.9	59	
sperm recombinant chromosome selection (+ simple embryo selection, 10 embryos)	nondestructive DNA sensing, ~8 loci chromosome transplantation sperm DNA must be labellable sperm DNA must be mechanically robust enough	.7 + [1 + (M(donors)/√2)] × (22/√46)	1 donor	3.9	23	i.e. 1 father
			2 donors	5.2	31	
			4 donors	6.3	38	
			10 donors	7.5	45	
			23 donors	8.4	50	
			300 donors	10.5	63	
			20,000 donors	13.1	79	
....with replacement of k chromosomes		simulated	1 donor, 2 chr	4.6	28	
		donors here are *additional* to the main parent contributing the rest of the haploid	1 donor, 5 chr	5.2	31	
			4 donors, 2 chr	4.9	29	
			4 donors, 5 chr	5.7	34	
			23 donors, 1 chr	4.6	28	
			23 donors, 2 chr	5.1	31	
			23 donors, 4 chr	5.9	35	
			23 donors, 10 chr	7.7	46	
haploid iterated editing stem cells	spermatogonial stem cell induction/culturing, transplantation, artificial spermatogonia aren't outcompeted in vivo, GFP insert activated in sperm OR in vitro oogenesis OR in vitro spermatogenesis	4 + (SDs × 15 / 2)	15 edits		8	0.56
		from kman. 5-15 edits / round, perhaps. haploid (1/2 effect) + 4 points for embryo selection (monoclonal edited SSCs).	50 edits		13	1.26
			100 edits		18	1.81
			200 edits		23	2.49
			300 edits		26	2.96
			500 edits		31	3.63
			1000 edits		39	4.66
ESC chromosome selection	same requirements as gamete chr selection + weak maintenance + ESC SCNT; but DNA more accessible	same as sperm+egg recomb. chromosome selection; see below	1 couple	6.5	39	
sperm and egg recombinant chromosome selection	nondestructive DNA sensing, ~8 loci chromosome transplantation sperm + egg DNA must be labellable sperm + egg DNA must be mechanically robust enough	[1 + (M(donors)/√2)] × (44/√46)	1 donor each	6.5	39	yes really! i think.
			2 donors each	9.1	55	
			4 donors each	11.2	67	
			10 donors each	13.5	81	
		"1 donor each" means one progenitor of each haploid, i.e. a single couple	23 donors each	15.3	92	
			300 donors each	19.7	118	do not do this
			20,000 donors each	24.9	149	do not do this
....with replacement of k chromosomes		simulated	1 donor, 2 chr	7.9	47	
		donors here are *additional* to the main parent contributing the rest of the haploid	1 donor, 5 chr	9	54	
			4 donors, 2 chr	8.4	50	
			4 donors, 5 chr	10	60	
			23 donors, 1 chr	7.8	47	first chromosome replacement gives +8 IQ points
		replacement and chr selection for each haploid separately	23 donors, 2 chr	8.8	53	
			23 donors, 4 chr	10.4	62	
			23 donors, 10 chr	14	84	
in vivo spermatogonial editing	delivery, low off-target rates, good efficiency, fluorescent tagging that shows in sperm; edited spermatogonia have to stay alive	6 + (SDs × 15 / 2) from kman. haploid (1/2 effect) + 6 IQ for full embryo selection	~50 edits?		15	1.26
iterated editing ESCs	maintain naive ESC epigenetic state, including imprinting, for months; ESC SCNT if imprinting not perfectly maintained, epigenetic editing OR full diploid epigenomic correction	6 + SDs × 15	15 edits		14	0.56
		from kman. 5-15 edits / round, perhaps. diploid (full effect)	50 edits		25	1.26
			100 edits		33	1.81
			200 edits		43	2.49
			300 edits		50	2.96
			500 edits		60	3.63
			1000 edits		76	4.66
haploid chromosome selection	nondestructive DNA sensing, 1 locus chromosome transplantation spermatogonial stem cell induction, transplantation, artificial spermatogonia aren't outcompeted in vivo, GFP insert activated in sperm OR in vitro oogenesis OR in vitro spermatogenesis	.7 + M(2 × donors) × (22/√46)	1 donor	2.5	15	
			2 donors	4	24	
			4 donors	5.3	32	
			10 donors	6.8	41	
		(haploid)	23 donors	7.9	47	
			300 donors	10.7	64	
			20,000 donors	14.2	85	
iterated meiotic selection	direct induction of meiosis full diploid epigenomic correction (OR one of IVO / IVS / iSSC transplantation, for half the effect)	M(2 × couples) × √rounds × (44/√46)	1 round, 1 couple	3.6	22	
			1 round, 2 couples	6.7	40	
			1 round, 4 couples	9.2	55	
			1 round, 10 couples	12.1	73	
		hard to determine bounds. estimates here assume rigid segments, but allow a round of chromosome selection at the end. without that assumption these numbers are optimistic by a lot, but probably still give a good general sense.	2 rounds, 1 couple	5.1	31	
			2 rounds, 2 couples	9.4	56	
			2 rounds, 4 couples	13	78	
			2 rounds, 10 couples	17.2	103	do not do this
			3 rounds, 1 couple	6.3	38	
			3 rounds, 2 couples	11.6	70	
			3 rounds, 4 couples	16	96	
			3 rounds, 10 couples	21	126	do not do this
			4 rounds, 1 couple	7.3	44	
			4 rounds, 2 couples	13.4	80	
			4 rounds, 4 couples	18.4	110	do not do this
			4 rounds, 10 couples	24.3	146	do not do this
			5 rounds, 1 couple	8.1	49	
			6 rounds, 1 couple	8.9	53	
			7 rounds, 1 couple	9.6	58	
			8 rounds, 1 couple	10.3	62	

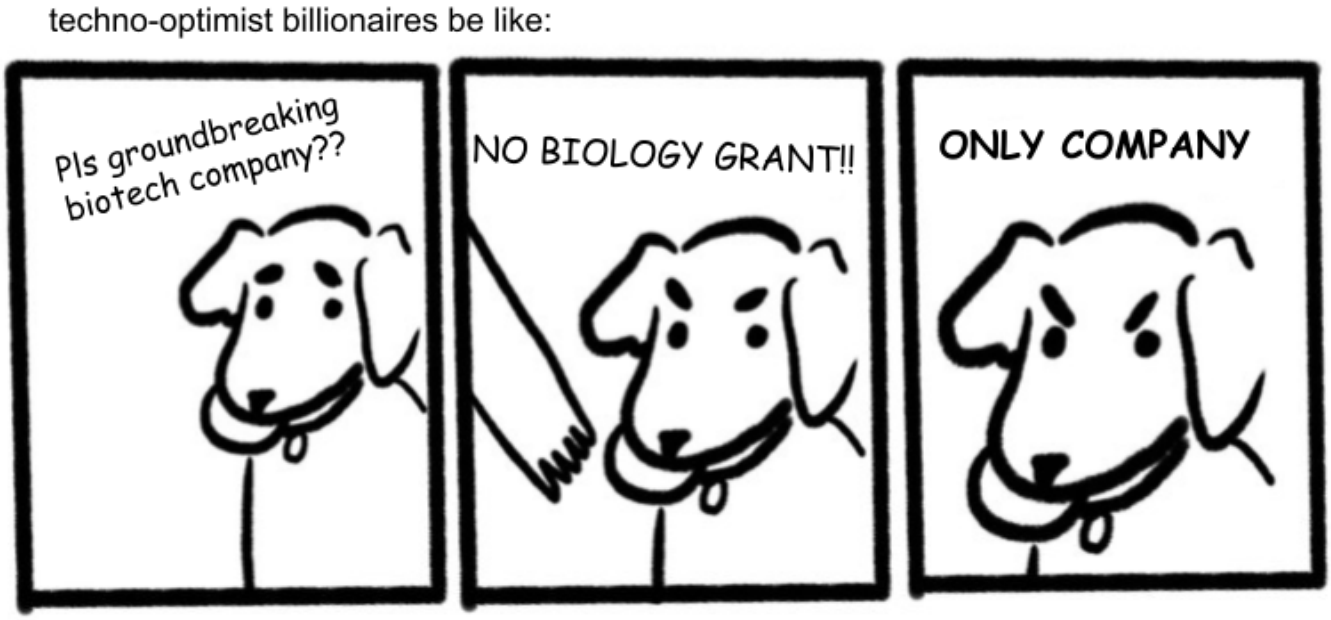
The list of specific methods that this table summarizes starts after the section Summary of genomic vectoring methods. This article is roughly organized from general to specific, first discussing things that apply to the whole area, and then later discussing specific methods.

I won't lie, this ~~article~~ book is a bit of a slog. You try writing a book about the state of the art of realistic germline engineering in a way that is automatically fun, and then get back to me, ok? But listen: There's lots of pictures, and some of them are good pictures. You could just skim through and look at the pictures, and then read more if there's something you really want to know. If you're looking for mathy stuff, or you just want to hear about the genomic vectoring methods, start here. Also, if you have some sort of oomph you might put into this (such as brainpower or moneypower), or if you're up for a recorded and published conversation, you can message me or email me, and I'll talk you through this stuff. Gmail address: berkeleygenomicsproject

Takeaways

- The bottleneck to strong human germline genomic engineering is not polygenic scores. The bottleneck also isn't really about making a cell with the genome you want. Rather, the bottleneck is **epigenomic correction**—making an engineered cell have the right epigenomic state so that it can make a healthy baby.
- **You would be surprised at how soon strong germline engineering can be made technically feasible.**
  - The scientific and technological precursors to strong human germline genomic engineering are ripening. There are several paths that look fairly feasible.
  - Peering through the looking glass at biotechnology, small feats may appear much easier than they really are, but big feats may appear much harder than they really are. It would probably take much less than a Manhattan project to make germline engineering real; the question is more about deciding to do it, rather than whether it can be done.
- This is an "enchanted field".
  - **Many advances could be combined with each other to make the methods safer, more powerful, and less expensive.**
  - That means that efforts are helpful not just in expected value, but with substantial probability. To some extent this is true of science in general. But in germline genomic engineering, different last-mile startups in the area will likely literally synergize, e.g. through licensing, collaboration, or serving as multiple inputs to other groups (see below).
  - Many methods have a nice innovation ramp. You could perform a few edits at first, then more and more; or transplant just one chromosome, then several; or do iterated meiotic selection for one or two rounds, then several.
  - The innovation ramp, plus the potential for combinations, sets the field up for a sustained ramp-up of energy once it gets started in earnest.
- The component technologies—reproductive, genomic—need more funding, talent, and project organization. **There are lots of things to do, and the existing academic funding and commercial investment landscape is far from efficient with respect to accelerating germline engineering technology.** Experts and their projects need more funding to go faster and create more public goods for frontier assisted reproductive technologies. Some other projects that could be created de novo:
  - **Create a primate research center focused on studying frontier reproductive technologies**, e.g. verifying that they produce healthy offspring. Make an atlas of single-cell RNA sequencing data in primate embryonic development.
  - More epigenomic sequencing for human reproduction. E.g. characterize the natural range of variation in epigenomic state between single cells within one early-stage human embryo.
  - More talent and funding for ovarian follicle culture, which is a likely way forward for making oocytes in vitro.
  - Some genomic vectoring methods could likely be applied today or very soon in animals. Testing these methods would give valuable feedback about the techniques and about what happens when you genomically vector strongly according to a PGS in one shot.
  - Gather more human phenotype/genotype data, especially for personality traits.

- Develop methods for intact chromosome transplantation. (Three remarkable and little appreciated facts: First, recombinant chromosome selection is a quite strong genomic vectoring method. Second, average chromosome selection with many donors is also quite strong. Third, chromosome selection might largely bypass the epigenomic correction problem, which is a major bottleneck for strong genomic vectoring. It might be infeasible, but for reasons decorrelated from other approaches.)



- **Scientists, technologists, investors, grantmakers, parents, policymakers, other stakeholders**:
  - If you want to help the field along by coordinating, starting projects, or funding work, we'd love to hear from you in DMs or at this gmail address: berkeleygenomicsproject
  - Consider joining us at the Reproductive Frontiers Summit 2025.

Disclaimers

I'm not a biologist. Parts of this article attempt to summarize some aspects of a large, complex, changing subarea of biology—the cell biology of human reproduction—as it relates to the possible future technology of human germline engineering. These summaries will necessarily be very incomplete, and unfortunately likely contain errors and confusions. Further, my knowledge is only a few steps from the shallow end, so for example I might say things based on outdated consensus. My hope is that the summaries, while lossy, will help others think about the subarea by pulling lots of threads together into one place and analyzing the basics of how those threads interact with germline genomic engineering.

This article is not intended as an explainer, but rather an attempt to summarize the state of the art for people interested in understanding or contributing to future developments. Unfortunately several of the sections are not independent from the rest of the article, so you might have to jump around and also internet search things.

In trying to find ways to implement strong human germline engineering, my understanding of what's possible, feasible, or easy seems to continually change, even aside from the fact that the field is progressing. There is a "layman's optimism" I've encountered in myself and in others. For example, to my layman's eyes, chromosome selection seems so simple—you just, you know, move the little guys around a bit until they're all together. But any specific plan has big holes in it (micromanipulator? nah, chromosomes are tiny; FACS? that'll probably break the DNA; do it on sperm DNA? it's highly compacted and inaccessible; etc.). Yes, a woman could extract some of her ovary tissue and then grow a myriad of eggs… if she knows someone who knows how to do ovarian culture well, and if ovarian follicle dominance doesn't get in the way too much. Yes, you could edit spermatogonial stem cells and transplant them into testes… but they'd be outcompeted by unedited spermatogonia and die out. I've tried fairly carefully to not overstate the feasibility of methods.

The math about the genomic vectoring power of different methods should be fairly solid, given the assumptions I make. The assumptions that connect the math with actual cells and DNA moving around are shakier; I make many simplifications, some knowingly and some not. The conclusions I give are based on simulations

of a simplified abstract model of genomic selection protocols, and is not based on using real DNA sequences of real genomes and operating on those. But I think that the qualitative conclusions should hold fairly well in most cases—e.g. comparisons of strength between different methods, and general ballpark estimates of strength. My hope is to communicate not "this is exactly how powerful these methods are" but rather "a natural first-order estimate says this method is really strong". I think history shows that "surprisingly large first-order estimate plus lots of complications" is, while very far from a sure bet, nevertheless often a very good thing to bet on.

In this article I speak solely on my own behalf.

The entire process of reproductive genomic vectoring (i.e. human germline genomic engineering) is likely to be complicated, especially at first, and hard to foresee. This article is not trying to address everything that would be required for reproductive GV. For example, any reproductive GV protocol must involve several health verifications, e.g. genome sequencing, epigenome sequencing, and morphological normality. This article just addresses the two core elements: genomic vectoring and epigenomic correctness.

Terms

Non-standard terms

Most terms in this article are standard. Some non-standard terms:

- Genomic vectoring (GV)
  - Unfortunately this conflicts with "germinal vesicle oocyte". In this article, GV always means genomic vectoring.
- Reproductive genomic vectoring
- Epigenomic correctness (EC)
- Chromosome selection
- Iterated meiotic selection (IMS)

Standard terms

Some standard terms with abbreviations:

- PGS (polygenic score)
- SNP (single nucleotide polymorphism)
- IVG (in vitro gametogenesis)
- IVO (in vitro oogenesis)
- IVS (in vitro spermatogenesis)
- ESC (embryonic stem cell)
- iPSC (induced pluripotent stem cell)
- PGC (primordial germ cell)
- SSC (spermatogonial stem cell)
- Combinatorial abbreviations:
  - Before an acronym: i means induced, h means human, m means mouse
  - After an acronym: LC means "-like cell"
  - So an miPSC is a mouse induced pluripotent stem cell, and an hiPGCLC is a human induced (primordial germ cell)-like cell .
- This is quasi-standard, but to clarify: I use -etic (genetic, epigenetic) to refer to a small number of DNA loci, and -omic (genomic, epigenomic) to refer to genome-wide effects.

Reproductive genomic vectoring

Reproductive genomic vectoring means making a baby who has a genome that was intentionally influenced, rather than solely by the natural reproductive process.

There are many downside risks, both technical and social, to reproductive genomic vectoring. See "Potential perils of germline genomic engineering". I'll address these elsewhere, along with the case in favor of germline engineering.

Genomic vectoring (GV)

Genomic vectoring (GV) means making a cell that contains a genome that has been modified to score highly according to some criterion. The genome could be diploid—a full complement of 46 chromosomes, two of each index, like all non-germline cells in your body; or it could be haploid—23 chromosomes, one of each index, like sperm or eggs.

Once you've genomically vectored a cell, you aren't done. Your GVed cell might just be a generic stem cell. It's not automatic that you can make a healthy baby from any old stem cell.

The epigenomic correctness (EC) problem is the problem of making cells that are epigenomically developmentally competent: they have the right epigenomic states so that they can contribute to growing a healthy baby. Any GV method requires some way of handling the epigenomic correctness problem. See the later sections "Reproductive GV and epigenomic correctness (EC)" and "Methods to handle epigenomic correctness".

As a view from 30,000 feet, the elements of a method for human germline engineering are to *know* and *make* the target *genome* and *epigenome* in a cell:

	know what your target looks like	make a cell with the target
genome	**polygenic scores (PGS)** constructed by genome-wide association studies and other analysis. requires phenotype+genotype data. clinicians inform parents of the meaning of available PGSes.	**genomic vectoring (GV)** editing or selection. selection methods: embryo selection, gamete selection, chromosome selection, iterated meiotic selection.
epigenome	**epigenomic correctness** epigenomic atlases of natural reproductive DNA (gametes, embryos). bisulfite seq, scRNA seq, scATAC seq, nanopore, ChIP-seq.	**epigenomic correction** in vitro EC-making; using natural reproductive DNA; in vivo reprogramming.

Berkeley Genomics Project
2025 license: CC BY-SA 4.0

Strong GV and why it matters

By "strong GV", I mean something a little nebulous. The term is meant to include methods that can greatly decrease the risk of several diseases in a future child, or substantially enhance some capacity, e.g. afford 30 IQ points or more. The term is meant to exclude simple embryo selection on normal numbers of embryos, and is meant to exclude editing a few loci.

Strength matters for reproductive genomic vectoring:

- A stronger GV method can be used to improve the baby's health on more dimensions (decreasing the risk of more diseases). Besides the obvious, decreasing disease risk is good because it helps to protect the baby against potential health problems introduced by the GV method itself.
- Strong GV removes many difficult tradeoffs between traits. With weak GV, parents have to evaluate whether they prefer to decrease their child's risk of diabetes by an additional 0.5% or to increase their child's expected IQ by a couple points (or something). With strong GV, the question becomes about what genomic foundation do the parents view as most desirable to give their children. Strong GV methods also remove tradeoffs between GV strength and the similarity of the resulting children.
- GV strength is generally interchangeable with cutting costs. Cutting costs is crucial for making GV technology widely available, in order to have the greatest benefit and to prevent problems with inequality.
- Strong GV would enable parents to have children with more opportunity for true genius of some flavor—scientific, scholarly, artistic, philosophical, political, communicative, technological, organizational. This would give the next generation high capability for both personal thriving as well as

intellectual and altruistic contributions, such as helping humanity navigate the rapidly changing world.

- If a strong GV can be made safe and soon, then we can more quickly demonstrate the large benefits of the technology. This will more quickly bring the scaled-up accessibility of very beneficial implementations of the technology.

Polygenic scores

Generally, the criterion for vectoring is a polygenic score—a function that predicts a trait from a genome. Genomic vectoring (GV) seeks to make a cell that scores highly according to the PGS, i.e. is predicted to make a baby who has a high degree of the trait/s.

The analyses in this article assume linear PGSes, i.e. PGSes that don't model gene-gene interactions. In practice, the criterion would be a PGS formed as a weighted sum of PGSes for various traits (or maybe a more complicated function of PGSes, such as an intersection of acceptable ranges). This article doesn't discuss what PGSes to use, how to use them, how to get them, how good they are, and what exactly they do and do not mean. In abstracting from those details, this article also doesn't account for the failure of PGSes to fully transfer between ancestry groups; it's worth keeping in mind that more data will be needed to afford the full opportunities of germline engineering to everyone.

IQ is used as a practical touchpoint to understand the power of different GV methods. It's convenient because it's a highly polygenic trait, and there already exist PGSes that *correlate* with IQ by at least 0.4; and we have some coarse sense of what the trait means. Also, it's a trait that to me seems quite important, in that I want to have clever kids, and I want other parents who want to have clever kids to be able to do so. That said, there are lots of other quite important traits that would be great to have PGSes for, from disease risks and longevity, to personality traits and more specific cognitive capacities (including ones that are also very relevant to intellectual achievement, e.g. courage and curiosity). Also, *I have not evaluated the literature on IQ PGSes.* If restricting to a PGS that only uses correlations that are actually causal would give a correlation less than .4, then all the estimates in this article about IQ specifically would have to be adjusted down proportionally (just multiply the raw SDs by the true correlation rather than by .4; or multiply the IQ points by (correlation/0.4)).

See "Embryo Selection For Intelligence", Branwen 2016[1], for more about the genetics of intelligence (and early discussion of GV methods).

Editing and selection

There are two broad types of methods of genomic vectoring (GV):

- *Editing.*
  - Basically, this means going into the DNA and deleting, adding, or replacing some of the DNA sequence.
  - Examples: double-strand break editing; base editing; prime editing; Fanzor editing.
  - In more words: editing involves (somehow, e.g. with a viral vector) delivering some molecules (such as CRISPR-Cas9 or another CRISPR system) into a cell. The molecules then cut or nick some nucleotides from the DNA, and/or chemically modify the DNA. Then, some ambient DNA repair machinery makes some other nucleotides replace the ones that were removed or mismatched, maybe with the help of another part of the editor. Thus the DNA gets changed at that location.
- *Selection.*
  - Basically, this means taking DNA segments that already exist in cells, and then assembling those segments into new combinations in a new cell.
  - Examples: (simple) embryo selection; iterated embryo selection; iterated meiotic selection; chromosome selection.
  - In more words: selection involves shuffling around chromosomes or large segments of chromosomes. The shuffling methods can be either natural, such as gametogenesis or fertilization, or artificial, such as cell fusion or mechanical chromosome manipulation. By shuffling DNA and using DNA sensing, cells with targeted DNA segments are created and identified.

[1] Gwern Branwen. 'Embryo Selection For Intelligence', 22 January 2016. https://gwern.net/embryo-selection.

Other types of genomic vectoring (GV)

Cloning is an edge case in that it targets a genome that already exists. This article doesn't discuss cloning because it has its own ethical problems, and it's not what most people want to do. Also, cloning is generally risky: most cells accumulate genetic damage, and so are dangerous to use to make a baby. Thus cloning would require another genomic vectoring method (as well as major epigenomic correction) to work safely.

Whole genome synthesis is excluded because I don't know about it. Reliably synthesizing tens of millions of base pairs, as would be required to make human chromosomes, is not currently feasible. It also may have disqualifying dual-use risks. If whole genome synthesis becomes feasible, then most of the pathways here are irrelevant. However, full epigenomic correction would be necessary. So, reproductive GV methods that somewhat bypass the epigenomic correction problem would still be front-runners until full epigenomic correction is solved.

Comparing editing and selection

Some broad remarks:

Editing produces out-of-distribution DNA

- Compared to selection, editing takes the resulting DNA more out-of-distribution.
  - Selection just stitches together large preexisting segments of DNA. Further, most methods do the stitching in places where natural reproduction might also have done the stitching (by recombination in natural gametogenesis). We can therefore be confident, at the local level of haplotypes (shortish segments of DNA on one chromosome), that genomic selection doesn't produce anything weird or dangerous.
  - Editing, on the other hand, cuts DNA at artificial spots. The most efficient editors, base editors, target single nucleotides. These editors thus create novel haplotypes. It's therefore harder to be confident that the results are safe. For example, if you edit a SNP, it could be that there are nearby rare variants that are correlated with the SNP.
  - There's at least one theoretical reason for worry: nearby base pairs are in strong linkage disequilibrium, and therefore might be coevolved with each other. If so, it's bad to modify a single base pair a in haplotype A to be the base pair b from haplotype B: some variants in haplotype A might be good or fine with the A allele a, but bad with the B allele b. In other words, the edit $A_0aA_1 \to A_0bA_1$ pairs b with A_0 and A_1, which it didn't coevolve with. There would have been very little selection pressure to get rid of a-associated variants in A_0 and A_1 that are bad when paired with b, since they don't in fact naturally appear with b. (I have heard there may be good reason to think this isn't a problem, but haven't followed up on this. For one thing, the number of such oddities would be quite small in the scheme of things, though also selected for being in important DNA regions.)

Editing is already minimally working technology

- Compared to most selection methods, editing is much more worked-out.
  - Except for simple embryo selection, any genomic selection will involve methods that aren't currently known, such as in vitro gametogenesis, in vitro meiosis, or intact chromosome transplantation.
  - Genetic editing, on the other hand, is a big industry with rapid progress and lots of existing methods that already work well. There will still be many details to work out, especially with iterated multiplex editing, and there are known and unknown variables that will make iterated editing difficult, but the basic principle is proven.
  - However, all strong genome editing methods require some form of epigenomic correction and/or maintenance.

Headwind of mutation

- There's always a headwind of mutation in GV (genomic vectoring) methods.
  - Most GV methods, of either kind, involve culturing cells for several generations. Many GV methods also involve doing other operations to cells, such as isolating and passaging single cells to another petri dish, or filtering cells according to some reporter.

  - These operations introduce risks of genomic degradation. Mutations are introduced at some slow but steady rate; survival in culture might tend to select for certain mutations; and passaging might further select for mutations.
    * For example, a cell being non-sticky is both cancer-associated and also makes it easier to isolate for monoclonal passaging. In lab settings, many stem cell lines carry oncogenic mutations.
  - This issue can probably be overcome by dovetailing whole genome sequencing and GV, to continually filter out cells with damaged DNA, but it's a pervasive complicating factor.
  - Some methods circumvent the headwind of mutation. Simple embryo selection and sperm chromosome selection probably don't interrupt the natural reproductive process enough to add much mutation.

Fixing de novo mutations

- Both editing and selection should be able to fix de novo mutations coming from the parents.
  - Organisms accumulate genetic mutations in their cells as they age; in particular, stored oocytes accumulate mutations (though at a reduced rate) and spermatogonial stem cells accumulate mutations.
  - Further, meiosis itself can cause mutations.
  - Editing and selection should both be able to fix these. Editing can just undo the change, at least for SNVs or other small changes. Selection can select the undamaged allele. (But selection can't fix damage in a Y-linked locus in male genomes, unless the damage only occurs in some of the Y chromosomes available.)
  - I'm not sure which is more efficient though. On the one hand, there may be around a hundred de novos in a given diploid genome assembled from the parents, so selection has a lot to work with. Also, selection can target larger kinds of damage without much additional cost. On the other hand, to fix a specific error, selection has to make a decision about a whole large segment (chromosome or significant fraction of a chromosome), so it seems like more selection power is expended in that toy example.

Bypassing natural integrity checks

- Artificial reproduction might let through more bad mutations and aberrant epigenomic states compared to natural reproduction.
  - Gametogenesis and fertilization involve many millions of cells undergoing various selective filters, e.g.: ability of gametogonia to proliferate, ability to appropriately respond to regulatory signals, completing recombination and the rest of meiosis without tripping too many DNA damage detectors, oocyte dominance contests during ovulation, physically passing through reproductive organs, and supporting early embryonic development. Germline cells with bad enough genomic or epigenomic problems don't naturally make it into a conceptus.
  - Artificial reproduction would skip some of these steps, removing some of the selection pressure towards integrity. The germline-like cells that would have died out in a harsher in vivo context, but make a conceptus artificially, might have more genomic or epigenomic problems.

Selection makes similar children

- By default, selection GV methods produce children with more genetic overlap than normal siblings.
  - Most of the discussion of GV methods in this article focuses on making a single child.
  - For people who want to have several children with reproductive GV, if they use a GV method that is a selection method, there would be a tradeoff. If they select fairly strongly according to some PGS, they're likely to make children who are significantly more similar to each other than normal siblings. Selection methods move whole large DNA segments around. If some DNA is anywhere nearby a patch of PGS-high-scoring alleles (like, in the same quarter of a chromosome or something), that DNA will tend to make it into many of the GV-selected cells. So GV-selected cells will tend to share that DNA more often than a random coinflip, as with natural random reproduction.
  - This does not apply to editing methods, which change a tiny fraction of DNA (less than one base pair in a million), out of mostly random natural DNA. (However, it's conceivable that some noticeable phenotypic similarity would be induced. For example, if multiple children received all

the same high-effect IQ-increasing edits, those edits might have had some specific flavor of effect on cognition, e.g. on personality.)

- I think most of the conclusions about selection methods wouldn't be qualitatively changed too much if the similarity / selection power tradeoff were taken into account. But, likely the GV power would decrease by some amount. The situation is complicated so it's hard to tell exactly how much; there is a lot of variation available in the genome to select from, even just in one person's genome. In particular, strong forms of selection GV would be able to alleviate the tradeoff, by simply promoting different high-scoring segments in different children.
- A simple touchpoint:
  * Suppose that a couple wants to have two children using a selection GV method. Suppose also that they want the children to have strictly half genomic overlap, as would two normal (non-twin) children. How much selection power can they get?
  * Suppose further for simplicity that we're only focusing on chromosome selection, and we only use the parents's chromosomes. In this case, the answer is: they can get at least half as much selection power as if they just made twins with all the highest-scoring chromosomes.
  * The strategy is this: Pick half of the chromosome indices 1—23. For that half, you pick the higher-scoring chromosome from each of the two parents. Both kids will receive a copy of those chromosomes. For the other half of indices, for each parent, you'll randomly pick which kid gets one homolog and which gets the other from that parent. You've used perfect anti-correlation in the latter half of the genome. Thus you concentrate the difference-between-kids chromosomes, and separate them out from the high-scoring-chromosomes.
  * You can do better than this by choosing the indices to select on separately for each parent, specifically to find where the score differences between the chromosomes are largest.
  * If the parents prefer 75% similarity instead of 50%, they can have 75% of the selection power with the same strategy—overlap on 75% of chromosomes, anti-correlate on the other 25%.
- Many of the estimates in this article about the power of chromosome selection ignore the 23rd chromosome, i.e. we assume we don't select over the 23rd chromosome. Partly that's just for modeling convenience, since the sex chromosome behaves differently than the other chromosomes. But also, it seems strange to imagine parents wanting to have the sex of all their children be the same, just to get a tiny bit more selection power. The difference is less than 5%; see the appendix "Ignoring the 23rd chromosome".

Selection and editing benefit differently from stronger PGSes

- As the strength of PGSes for target traits increases, selection will directly proportionally increase in power; but editing will have a different curve of increase in power.
  - The issue for editing is that it is bottlenecked on number of edits. For a given combined PGS, you pick the top k edits to make, where k is however many edits you expect to be able to make. Improving one of the PGSes marginally may find lots of weak-effect variants without finding any or many large-effect variants, which wouldn't increase the power of a k-edit protocol.
  - Selection, on the other hand, takes advantage of all the information across the whole genome at the same time, so to speak. If you simultaneously improve the correlation all of your component PGSes with their target traits by a factor of f, you then also multiply by f the correlation of your combined PGS with the combined target trait. Selection GV methods, without changing the protocol power, would then have their final effects multiplied by f as well.
  - The same applies for adding more PGSes (that is, mixing in a newly constructed PGS for a previously untargeted trait). Both editing and selection will get more powerful (in terms of total effects on the new combined score), but selection will improve much more.
- However, the scale situation is not necessarily so bad for editing.
  - In some regime, including right now, there's significant uncertainty about which exact SNP within a segment of 5ish SNPs is the actual causal SNP[2]. Improved PGSes should narrow that down, which should substantially improve the expected effects. Selection would benefit little or not at all from more precise causal information (at least within the same population as the source of the PGS): selection already captures causal variants by selecting for larger DNA segments, which contain the causal variant and its nearby correlates.
  - For many/most disease traits, the gains for any method are rapidly diminishing. The first few changes in disease alleles, starting with a normal genome, will make much more difference in

[2]GeneSmith, and kman. 'How to Make Superbabies', 19 February 2025. https://www.lesswrong.com/posts/DfrSZaf3JC8vJdbZL/how-to-make-superbabies.

absolute disease probability, compared to changes made to an already especially low-risk genome. E.g. a move from 3 SD to 5 SD low risk for some disease might represent only a .0001 difference in disease probability, but a move from -1 SD to 1 SD might be a .05 difference, or something. (In contrast, returns for IQ variants don't diminish nearly as quickly.)
  - There may be many rare, high-effect variants, as is the case for some traits. Improved PGSes would progressively discover these, thus unlocking marginally higher-effect edits. However, these rare variants would plausibly come with an increased risk of unintended effects, especially if they have a strong enough impact on the target trait that they make it into the top k over more common variants.

Selection distributes Euclideanly, editing distributes Hammingly

- Selection distributes SDs of selection power Euclideanly to component PGSes.
  - See https://tsvibt.blogspot.com/2022/08/the-power-of-selection.html#6-selecting-for-multiple-scores.
  - This means that if you can select a genome so that it is k standard deviations (SDs) extreme, you get k^2 points. Then you distribute those k^2 points to different traits (for which you have PGSes). If you give n points to a trait, that trait will be $r\sqrt{n}$ SDs extreme, where r is the correlation coefficient of the PGS with the trait.
  - For example, say you select a genome to be 5 SDs extreme. You get 25 points. You could allot them all to IQ, and then you'd get $.4 \times \sqrt{25} = 2$ SDs $= 30$ IQ points, because we have a .4 PGS for IQ. Or you could put 16 points into IQ, and use up the rest of your points by putting 3 points into each of 3 health traits. Then you'd get $.4 \times \sqrt{16} = 1.6$ SDs $= 24$ IQ points, and $.3 \times \sqrt{3} \approx .5$ SDs on each of the health traits (if, say, they each have a .3 PGS).
- Editing distributes editing power (quasi-)linearly.
  - Editing moves the genome around in genome-space, which is a sort of Hamming space. At a given level of development, an editing protocol can take some number of steps—i.e. make some number of edits.
  - With editing, the question is: what are the most trait-positive variants that you know about, how many aren't already in the genome you're editing, and how many edits can you make given the mechanics of cell culture and editing molecules.
  - Given a number of edits and multiple target PGSes, you have to simply apportion each edit-slot to one of PGSes. This is "linear" distribution.
  - (However, for each trait, some edits will have a larger SD effect on that trait than others. So the literal total number of SDs (which is not an important metric) is not fixed even with a fixed edit count; hence "quasi-linear".)

Selection beats editing, then editing beats selection

- Selection is stronger than editing when both are weak; editing is stronger when both are very strong.
  - In other words: A weak version of selection will have a greater effect on traits compared to a weak version of editing. On the other hand, strong editing has a greater effect than strong selection.
  - Examples of weak/strong versions of selection:
    - chromosome selection on one/many donors,
    - iterated recombinant selection with few/many iterations,
    - simple embryo selection with an ordinary/gigantic number of eggs.
  - Weak editing is modifying up to a couple hundred loci; strong editing is editing many hundreds or thousands of loci. (In the limit, editing is as powerful as whole genome synthesis.)
  - For polygenic traits such as IQ and many health, longevity, and other cognitive traits, each genetic variant has a very small effect on the trait. For example, all or very nearly all IQ variants will have less than a .5 IQ point effect.
  - Weak selection is stronger than weak editing. Weak selection can weakly harness the power of the entire suite of available PGSes. Weak editing can only edit, say, the top few dozen most effective variants, and has to ignore all the others.
  - Strong editing is stronger than strong selection for two reasons. First, strong editing can add variants that aren't available in the parents's genomes, but that are known to be good because other people have them. Second, editing can, in principle, add variants that no human has, but that are believed to be good for some other reason (e.g., a variant coding region improves a protein's efficiency; or the variant is tested in chimps and improves something).
  - In any case, strong versions of either type of vectoring are sufficient to get enormous benefits.

Editing can reach more places than selection

- In general, editing can be applied to cells in more contexts than selection methods.
  - Editing can be done by delivering some molecules to some cells, wherever those cells are. For example, it's even possible to edit adult tissue in vivo, though with difficulty and with low efficiency. A practical example is that one might be able to edit gametogonia that inhabit gonadal tissue in vitro.
  - Most selection methods, on the other hand, require filtering cells and DNA. E.g. chromosome selection requires moving specific chromosomes between specific cells; iterated recombinant selection requires fusing specific haploid cells.

The ceiling of safe vectoring

- Some GV methods that are likely to be feasible would hit the upper limit of what would be safe.
  - Generally speaking, no trait should be pushed too far. Some examples:
    - IQ, if pushed too high, would produce severe physiological problems. See the bullet point "Traits outside the regime of adaptedness" here.
    - Health traits would likewise have unknown physiological effects if pushed too far.
    - Personality traits are likely bad at extremes. As a possible example: slightly low conscientiousness might produce creativity, and slightly high conscientiousness might produce diligence and carefulness; but very low conscientiousness might produce dysfunctional erraticness and psychopathy, and very high conscientiousness might produce neuroticism / OCD. Slightly low agreeableness might produce leaders and independent thinkers, and slightly high agreeableness might produce good team players and caregivers; but very low agreeableness might produce violent psychopathy[3], and very high agreeableness might produce exploitability or subservience.
  - Editing, if scaled up enough, would hit limits. Likewise chromosome selection with many parents and iterated meiotic selection would also hit limits.

Many GV methods synergize

Many methods for genomic vectoring (making a cell with a genome that's been nudged in some direction) can be combined with each other to increase power and/or decrease costs. There are many possible combinations of different methods, and until some GV technologies are ready to be applied it doesn't matter exactly how well various combinations work. To give some main examples:

- **Chromosome transplantation** would enhance most GV methods.
  - For most methods that operate on multiple cell lines, chromosome transplantation would allow gains in chromosomes from different cells to be aggregated.
  - E.g. if you iteratedly CRISPR edit some cells, you may have chromosomes in multiple different cells that successfully got the most target edits, and the least DNA damage. It would be nice to then just move those chromosomes together into one cell.
  - E.g. see the appendix "The costs of iterated meiotic selection" on how chromosome selection makes IMS less costly / more powerful by gathering together chromosomes that have target crossovers.
  - E.g. if editing isn't available, chromosome transplantation would enable preventing monogenic homozygous diseases by replacing one chromosome from a donor without the disease allele.
- **Editing** would enhance most GV methods.
  - For most methods that operate on culturable cell lines (so that you can sequence some cells to check for damage), editing would allow spot-correcting any damage, whether due to parental de novos, culturing, or abnormal cellular processes (e.g. inducing meiosis might cause DNA damage).
  - For selection methods, editing can introduce variants not present in the starting population of cells. So it can prevent homozygous disease alleles from being passed on, and it can enhance traits beyond what would be achievable just from the starting population, e.g. by editing in known helpful rare-ish variants.
- **Embryo selection** can be combined with any method that creates a viable haploid gamete-like cell, effectively selecting the complementary gamete as well.
- DNA sensing can help any selection method.

[3] Miller, Joshua D., and Donald R. Lynam. 'Understanding Psychopathy Using the Basic Elements of Personality'. Social and Personality Psychology Compass 9, no. 5 (2015): 223–37. https://doi.org/10.1111/spc3.12170.

GV synergy has two very auspicious consequences:

- If a GV method can be made to work, it's unlikely to be a dead end. Even if a method is delayed, can't produce much GV on its own, or is costly, there's a good chance it will still be useful as an admixture to another successful GV method. This somewhat decreases the risk a research project faces that its products will not be interesting or useful.
- Once any nontrivial GV method is working, it's likely that there could be a large, fast increase in the strength of the strongest combined GV method. As mentioned above, strength is important for protecting the baby from possible side-effects of the GV method, for making the technology widely available, and for granting greatly enhanced capacities.

GV synergy also highlights the key bottleneck in reproductive GV: dealing with the epigenomic correctness problem. There are several synergistic paths to GV which could work on their own or together, so I'm not very concerned that we won't have a good combined GV method. However, any full reproductive GV method requires solving the EC problem, whether with an EC-making (epigenomic correction) method or by bypassing the EC problem (e.g. using natural reproductive DNA).

Reproductive GV and epigenomic correctness (EC)

Reproductive genomic vectoring means making a cell with a vectored genome, and then making a *healthy* human baby from that cell.

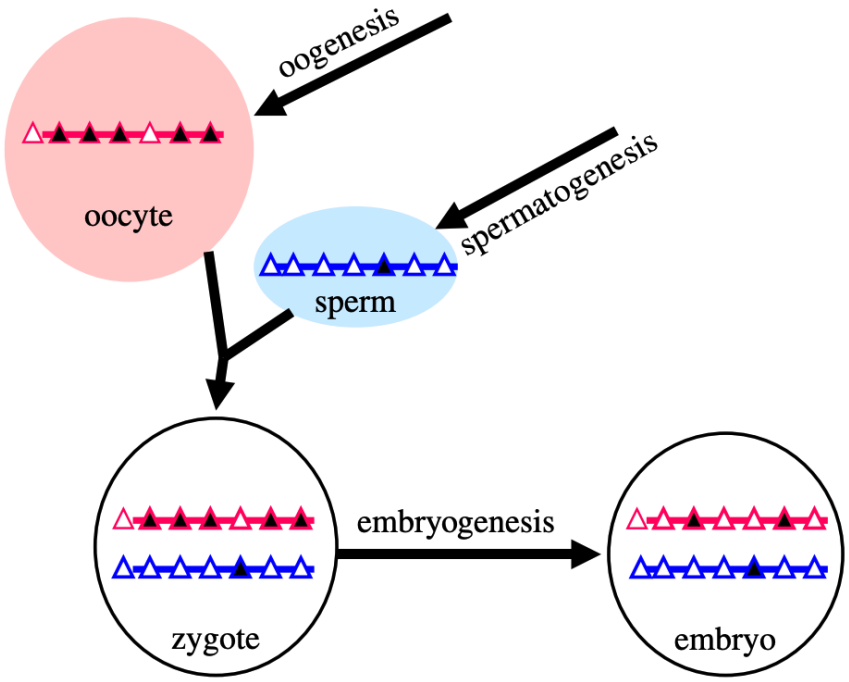
The epigenomic correctness problem

To reliably produce a healthy human baby from a cell, you must make sure that the cell has the right epigenomic state. The chemical markers on and around the DNA—e.g. CpG methylations and histone modifications—should look like the markers on natural reproductive DNA (eggs, sperm, zygotes).

Suppose you just take the nucleus from a somatic cell, or even a stem cell, and stick the nucleus into an oocyte and try to implant it. There's a high chance it won't grow, or if it does grow, there's a high chance it will miscarry or be stillborn, or if it isn't miscarried or stillborn, there's a high chance the resulting baby will be sick and deformed. That's because a random cell won't have the correct epigenomic state. This is the "epigenomic correctness" problem. (Or "epigenomic competence".)

In natural reproduction, gametogenesis produces gametes (sperm or eggs). The early stages of germline development leading up to gametogenesis involve wiping clean the epigenomic state of the stem cells that will become gametes. Genome-wide reprogramming happens in both sperm and eggs, but there are also sex-specific epigenetic imprints: spermatogenesis produces spermatozoa (sperm) with *paternal* epigenetic imprinting, and oogenesis produces oocytes (eggs) with *maternal* epigenetic imprinting. These sex-linked epigenomic imprints are different between the sperm and the egg, and are necessary for healthy growth of the fetus[4]. An illustration:

[4]Metacelsus. 'Epigenetics of the Mammalian Germline'. Substack newsletter. De Novo (blog), 23 December 2022. https://denovo.substack.com/p/epigenetics-of-the-mammalian-germline.



(Figure 1 from Kelsey and Feil (2013)[5].)

The triangles represent CpG islands in the genome; the black ones are methylated.

It's not known exactly what imprints are necessary or sufficient for healthy development. Some research has epigenomically sequenced human gametes and embryos[6], but we don't have a full picture (e.g., the cited paper doesn't look at histone modifications). It's probably something like several dozen to a couple hundred epigenetic marks across the sperm and oocyte genomes. (Over a hundred sex-linked differences are known, but many are controlled by others and therefore are not independently necessary.) **Knowing this information is crucial for most reproductive GV methods.** We'd also want detailed information on the epigenetics of fetal development so that we can tell if something is wrong. There hasn't yet been a definitive "gold standard" agreed on for what should count as demonstrating epigenetic safety for making human babies. Some of the main elements would be multi-omic similarity to natural gametes, safety demonstrations in animal models, and morphologically and multi-omically normal embryonic development up to 14 days.

There should be more funding for epigenomic sequencing of germline cells:

- Natural germline cells. Tissue from ovaries and testes from different points in the lifecycle, taken from non-human primates or from humans if ethical (e.g. using tissue from miscarried fetuses with informed consent of the parents). Primordial germ cells, gametogonia, gametocytes, gametes, as well as support tissue (gonadal cells such as granulosa cells and Sertoli cells).
- Especially, normal healthy gametes (sperm and eggs).
- Early embryonic tissue in non-human primates, to get a clearer picture of what healthy development looks like epigenomically.

Broadly speaking, methylations comprise most of the important sex-linked epigenetic imprinting marks, while histone modifications are at least to a large extent downstream of methylations. In sperm, most (around 90%) of the histones are replaced by protamines for tighter packing. However, at least some of the remaining sperm histones at paternally expressed genes retain modifications, so we definitely cannot rule out that some histone modifications are a necessary element of healthy paternal DNA[7]. Tanaka and Watanabe (2023) suggest that the main problem with ROSI is that round spermatids have too many histones that

[5]Kelsey, Gavin, and Robert Feil. 'New Insights into Establishment and Maintenance of DNA Methylation Imprints in Mammals'. Philosophical Transactions of the Royal Society B: Biological Sciences 368, no. 1609 (5 January 2013): 20110336. https://doi.org/10.1098/rstb.2011.0336.

[6]Jima, Dereje D., David A. Skaar, Antonio Planchart, Alison Motsinger-Reif, Sebnem E. Cevik, Sarah S. Park, Michael Cowley, et al. 'Genomic Map of Candidate Human Imprint Control Regions: The Imprintome'. Epigenetics 17, no. 13 (9 December 2022): 1920–43. https://doi.org/10.1080/15592294.2022.2091815.

[7]Ishihara, Teruhito, Oliver W. Griffith, Shunsuke Suzuki, and Marilyn B. Renfree. 'Presence of H3K4me3 on Paternally Expressed Genes of the Paternal Genome From Sperm to Implantation'. Frontiers in Cell and Developmental Biology 10 (10 March 2022). https://doi.org/10.3389/fcell.2022.838684.

haven't yet been replaced with protamines[8]. In trying to clone macaques, Liu et al. (2018) found that histone modifications were partly determinative of cloning success, so it is possible for histone modifications to matter a lot (though this effect could go away with otherwise normal methylations)[9].

It's not known how much loss of imprinting occurs naturally in somatic cells, or how much would occur as the result of operations involved in GV methods (such as inducing pluripotency with Yamanaka factors, inducing naivety with super-SOX, culturing in vitro, editing, or mechanical manipulation).

It's not known how to take a non-reproductive cell, and then correct its epigenomic state so that it can be a viable gamete or zygote (1-cell embryo). Such an epigenomic correction (EC-making) method would be one kind of solution to the EC problem.

Applying reproductive GV in humans would require a way to ensure epigenomic correctness with pretty high confidence, *before* making the first baby. However, there are several strong GV methods, such as iterated CRISPR editing or chromosome selection via whole cell fusion, that could be tried soon in plants, mice, or other animals.

Sex-linked epigenetic imprinting isn't the only aspect of the on-DNA epigenomic state that's needed for healthy development. In the early embryo, for example, the DNA has to be broadly demethylated like a natural naive ESC, so that embryonic cells can differentiate into all the tissues of the conceptus. Since broad DNA demethylation occurs naturally, given the environment provided by the egg, we may or may not have to worry about it specifically.

There are also other characteristics of sperm and eggs, or of zygotes, that are important for development, besides on-DNA epigenetic marks. These will be discussed briefly in the later sections on in vitro gametogenesis. The task of epigenomic correction is more narrow: we just need to get the DNA itself in a developmentally competent state, and then if necessary we can use a donor egg and/or sperm to give the rest of the needed support for fertilization and development.

The near-miss hazard of epigenomic correction

Do we actually need to have the right epigenomic state? Organic life is robust, it self-corrects. What if we just put the DNA we want into an oocyte (after getting rid of the DNA already there) and then tell it to grow?

That method is called SCNT—somatic cell nuclear transfer, a.k.a. cloning—and it's how the first cloned mammal was created. In 1997, Wilmut et al. announced they had made a lamb, number 6LL3, from an adult sheep mammary gland cell using SCNT[10]. (That very 6LL3 grew up to be none other than Dolly the sheep.)

[8]Tanaka, Atsushi, and Seiji Watanabe. 'How to Improve the Clinical Outcome of Round Spermatid Injection (ROSI) into the Oocyte: Correction of Epigenetic Abnormalities'. Reproductive Medicine and Biology 22, no. 1 (9 February 2023): e12503. https://doi.org/10.1002/rmb2.12503.

[9]Liu, Zhen, Yijun Cai, Yan Wang, Yanhong Nie, Chenchen Zhang, Yuting Xu, Xiaotong Zhang, et al. 'Cloning of Macaque Monkeys by Somatic Cell Nuclear Transfer'. Cell 172, no. 4 (8 February 2018): 881-887.e7. https://doi.org/10.1016/j.cell.2018.01.020.

[10]Schnleke, A E, K H S Campbell, Ian Wilmut, J McWhir, and AJ Kind. 'Viable Offspring Derived from Fetal and Adult Mammalian Cells', 1997, https://www.nature.com/articles/385810a0.pdf.



Figure 2 Lamb number 6LL3 derived from the mammary gland of a Finn Dorset ewe with the Scottish Blackface ewe which was the recipient.

(Figure 2 from Wilmut et al. 1997.)

Dolly the sheep had 276 siblings who didn't make it. Of the 277 mammary epithelium cells that Wilmut et al. did SCNT to, 29 grew to the blastocyst stage and were transferred into recipient ewes; only one was born. (They also transferred embryos derived by SCNT from embryonic stem cells and from fetal cells; those had a better survival rate, though still bad.) That should be alarming; something is at least quite unnatural about this process.

But maybe it's fine. More than 10% of normal pregnancies miscarry anyway; maybe, as along as a fetus grows for the full term, it'll be fine? Could we simply do SCNT with adult human cells? Perhaps after editing them? No. Don't do that.

Just because a fetus is born alive, does not mean it's healthy. There are plenty of possible developmental abnormalities that don't kill a fetus, but that do severely affect the health or lifespan of the offspring. Observations of 1000 cloned dogs show substantial fractions of them have significant developmental abnormalities such as cleft palate and muscular hypermyotrophy[11]. In cattle, the situation is not good: "On average, 42% of cloned calves died between delivery and 150 days of life; the most common abnormalities were: enlarged umbilical cord (37%), respiratory problems (19%), calves depressed/prolonged recumbency (20%) and contracted flexor tendons (21%)."[12]. Survival and deformity rates vary by species, by cell source, and by IVF method. But it's a minefield, and humans definitely do not have the *least* finicky mammalian reproductive epigenetics.

Part of the reason SCNT doesn't work to reliably make a healthy offspring is that sex-specific epigenetic imprinting is required for healthy development, even if it's not strictly required for minimal viable development[13]. The donor oocyte has enough material (e.g. transcription factors) in the cytoplasm that it can sometimes nudge the transferred DNA enough to sort of muddle through development. But many epigenetic marks will still be missing in cells from adults, having degraded during growth, and development is very prone to be abnormal. Optimizing your protocol to increase efficiency—to make more artificial zygotes make

[11]Olsson, P. Olof, Yeon Woo Jeong, Yeonik Jeong, Mina Kang, Gang Bae Park, Eunji Choi, Sun Kim, Mohammed Shamim Hossein, Young-Bum Son, and Woo Suk Hwang. 'Insights from One Thousand Cloned Dogs'. Scientific Reports 12, no. 1 (1 July 2022): 11209. https://doi.org/10.1038/s41598-022-15097-7.

[12]Smith, Lc, J Suzuki Jr, Ak Goff, F Filion, J Therrien, Bd Murphy, Hr Kohan-Ghadr, et al. 'Developmental and Epigenetic Anomalies in Cloned Cattle'. Reproduction in Domestic Animals 47, no. s4 (2012): 107–14. https://doi.org/10.1111/j.1439-0531.2012.02063.x.

[13]Metacelsus. 'Epigenetics of the Mammalian Germline'. Substack newsletter. De Novo (blog), 23 December 2022. https://denovo.substack.com/p/epigenetics-of-the-mammalian-germline.

it to a live birth—does not necessarily decrease your rate of deformities nearly enough.

That's the near-miss hazard of epigenomic correction. You correct the epigenomic state enough to make a live birth, but not enough for reliable healthy development. Not good.

I'm going to keep giving examples so you know I'm not making this up.

The most direct observation we can make about imprinting problems is uniparental disomy, explained to me by Tassilo Neubauer. In uniparental disomy, a child inherits two chromosomes of the same index from one parent. (In contrast to trisomy, the chromosome of that index from the other parent is not present for some reason.) Depending on the chromosome index, many people with disomy are fine, but many others have serious health issues. Some of this is due to homozyosity, if the chromosomes are identical; but some is due to the incorrect imprinting of the uniparental chromosomes (maternal-maternal or paternal-paternal). Read more at the UPD database[14] and in Neubauer's short review[15]. For a cell with not just one, but all chromosomes incorrectly imprinted, the effects would be lethal, but if not lethal then quite bad.

Nayernia et al. (2006) derived mouse haploid sperm-like cells from embryonic stem cells in vitro. Then they used those sperm-like cells to fertilize mouse oocytes and implant the resulting embryos. The ones that were born alive were over- or under-grown, and died 5 days to 5 months after birth. They confirmed that the experimental mice had abnormal methylation patterns in sex-linked imprinting regions, specific to whether they were over- or under-grown[16].

Recently, Li et al. (2025) made bi-paternal mice by making genetic edits to tweak the expression of regions that would normally be controlled by the correct imprinting. This worked, in that many abnormalities were prevented… but the mice that were born were sick, and almost all died[17]. Earlier, Li et al. (2016) had done a similar thing to make bi-maternal mice. The research process involved creating, along the way, severely undergrown bi-maternal mice[18].

Mitalipov et al. (2002) applied SCNT to rhesus macaques, making embryos from blastomeres—very early embryonic cells, which ought to be fairly epigenetically normal. Of 30 embryos transferred into 11 monkeys, only 1 pregnancy resulted. It grew to term, but was stillborn[19]. The authors guess the fetus died of asphyxiation by the umbilical cord, but one wonders if that was induced by some developmental problem caused by epigenetic abnormality; asphyxiation stillbirths are rare, and recall that cloned cattle frequently have enlarged umbilical cords[20]. Since the zygote genome is quickly reprogrammed (almost entirely demethylated) soon after fertilization, it stands to reason that there could be a difference in SCNTed ESCs that has significant consequences. For example, chromatin condensation of the oocyte genome, which would be messed up by ESC SCNT, might be important for early development[21].

Liu et al. (2018) cloned cynomolgus macaques with SCNT[22]. They used a treatment that altered the histone modifications in the DNA to be cloned. This greatly improved the success rate, and they made two apparently healthy offspring from fetal fibroblast cells. But they also tried the same method using somatic cells from an adult macaque. Of the few pregnancies that took, 2 miscarried late-term, and 2 were born

[14]Liehr, Thomas. 'Cases with Uniparental Disomy', 2025. https://cs-tl.de/DB/CA/UPD/0-Start.html.

[15]Morpheus. 'What Uniparental Disomy Tells Us About Improper Imprinting in Humans', 28 March 2025. https://www.lesswrong.com/posts/49Y2WDxStNBWMCBku/what-uniparental-disomy-tells-us-about-improper-imprinting.

[16]Nayernia, Karim, Jessica Nolte, Hans W. Michelmann, Jae Ho Lee, Kristina Rathsack, Nadja Drusenheimer, Arvind Dev, et al. 'In Vitro-Differentiated Embryonic Stem Cells Give Rise to Male Gametes That Can Generate Offspring Mice'. Developmental Cell 11, no. 1 (1 July 2006): 125–32. https://doi.org/10.1016/j.devcel.2006.05.010.

[17]Li, Zhi-kun, Li-bin Wang, Le-yun Wang, Xue-han Sun, Ze-hui Ren, Si-nan Ma, Yu-long Zhao, et al. 'Adult Bi-Paternal Offspring Generated through Direct Modification of Imprinted Genes in Mammals'. Cell Stem Cell, 28 January 2025. https://doi.org/10.1016/j.stem.2025.01.005.

[18]Li, Zhikun, Haifeng Wan, Guihai Feng, Leyun Wang, Zhengquan He, Yukai Wang, Xiu-Jie Wang, Wei Li, Qi Zhou, and Baoyang Hu. 'Birth of Fertile Bimaternal Offspring Following Intracytoplasmic Injection of Parthenogenetic Haploid Embryonic Stem Cells'. Cell Research 26, no. 1 (January 2016): 135–38. https://doi.org/10.1038/cr.2015.151.

[19]Mitalipov, Shoukhrat M., Richard R. Yeoman, Kevin D. Nusser, and Don P. Wolf. 'Rhesus Monkey Embryos Produced by Nuclear Transfer from Embryonic Blastomeres or Somatic Cells1'. Biology of Reproduction 66, no. 5 (1 May 2002): 1367–73. https://doi.org/10.1095/biolreprod66.5.1367.

[20]Smith, Lc, J Suzuki Jr, Ak Goff, F Filion, J Therrien, Bd Murphy, Hr Kohan-Ghadr, et al. 'Developmental and Epigenetic Anomalies in Cloned Cattle'. Reproduction in Domestic Animals 47, no. s4 (2012): 107–14. https://doi.org/10.1111/j.1439-0531.2012.02063.x.

[21]Luciano, Alberto Maria, and Valentina Lodde. 'Changes of Large-Scale Chromatin Configuration During Mammalian Oocyte Differentiation'. In Oogenesis, edited by Giovanni Coticchio, David F. Albertini, and Lucia De Santis, 93–108. London: Springer, 2013. https://doi.org/10.1007/978-0-85729-826-3_7.

[22]Liu, Zhen, Yijun Cai, Yan Wang, Yanhong Nie, Chenchen Zhang, Yuting Xu, Xiaotong Zhang, et al. 'Cloning of Macaque Monkeys by Somatic Cell Nuclear Transfer'. Cell 172, no. 4 (8 February 2018): 881-887.e7. https://doi.org/10.1016/j.cell.2018.01.020.

alive. Quoting:

> Infant A showed normal head circumference but impaired body development at birth and died 3 hr later due to apparent respiratory failure. Infant B had apparent normal head and body development and showed normal breathing and food and water intake but died 30 hr later with respiratory failure (see Data S4).

In humans, epigenetic abnormalities have been somehow associated with serious disorders and cancer[23].

So that's why you can't just make a human baby without knowing what you're doing: You stand a high risk of making a baby with developmental abnormalities that weren't severe enough to abort the fetus, but are severe enough that the child is suffering. If for some reason the moral consequences of that aren't enough to dissuade you, consider that other people would ban you and your children and your children's children and your artificial children and any similar research for 1000 years.

Going from 0.01 to 0.99

Reproductive genomic vectoring involves composing multiple biotic processes together: gene editing and repair, mitosis and meiosis, folliculogenesis, gametogenesis, induction of stem cell states, embryogenesis.

Biotic processes—cell division, gametogenesis, embryogenesis, up- and down-regulation of genes, cell differentiation—are both robust and noisy.

They're robust, in that even if not every cellular process goes exactly according to the evolutionary design, the end result may be almost as good, or even exactly as good. Error correcting mechanisms such as DNA repair, methylation maintenance, self-perpetuating gene regulatory network states, and homeostatic feedback in general, can bring and keep a cell on track (probably, approximately), as long as the cell hasn't been too extremely perturbed.

Biotic processes are also noisy at every step. DNA copying is imperfect, chromosome synapsis is imperfect, DNA breakage repair is imperfect, and so on.

Because of the noise and the self-correction, surprising things happen. Embryonic stem cells cultured in a certain way seem to, in some small fraction, spontaneously undergo meiosis, despite the absence of most of the preconditions for normal gametogenesis. Somatic cells transplanted into enucleated oocytes will, in some small fraction, develop into healthy offspring. This means that with many challenges, we don't exactly start at 0, but rather we start at 0.01: We can do them, but only rarely, or perhaps slowly and at high cost, and with results whose quality is both poor and unknown.

In general, starting at 0.01 is not good enough:

- A main reason, as discussed in the previous section, is that for some processes, near-misses are very costly.
- Importantly, the development and safety validation of the technology involves lots of experimental iterations. If the inner loop of the experiments—the biotic processes that you have to recapitulate frequently and in great numbers—is slow and costly, you get much less rapid feedback. It's also harder to compose multiple methods together if they're unreliable, so it's harder to get especially informative end-to-end feedback. And it is harder to share and replicate methods.
- Also, ultimately we want the whole integrated protocol to be inexpensive and scalable as well as knowably, consistently safe. Inexpensive scalable methods would support bringing reproductive genomic vectoring to more people, and make a stronger case that reproductive GV won't increase inequality due to differential access.
- Errors may add up.
  - For example, a process for creating ovarian follicles might produce medium-quality primordial follicles, that sort of resemble natural primordial follicles. But then if you use the artificial follicles to mature some oocytes, you produce very low-quality oocyte-like cells, because the follicles can't really give the oocytes all the support they need.
- Poor efficiency ~~adds~~ multiplies up.

[23] Trasler, Jacquetta M. 'Gamete Imprinting: Setting Epigenetic Patterns for the next Generation'. Reproduction, Fertility and Development 18, no. 2 (2006): 63. https://www.publish.csiro.au/rd/Fulltext/RD05118.

- If you can make 1% of your iPSCs into fully competent mature oocytes, you can eat the cost of that low efficiency. But if only .5% of your immature oocytes become mature oocytes, and only .5% of your artificial mature oocytes can be fertilized and start growing, and only 1% of those embryos can be implanted, then your costs for the full process would be… *crunches numbers*… "really big".
- For some processes, poor efficiency can be compensated for by amplifying the desired cells at each stage, e.g. by filtering and proliferating cells. But this adds time and complexity costs, and doesn't work for everything (e.g. you can't straightforwardly proliferate oocytes).

So we can't be satisfied with 0.01. We have to get much closer to 1.

However, none of the above reasons demand getting to literally 1. Because natural biotic processes are noisy, our bar for quality does not have to be 100%. Babies born naturally have birth defects at a rate greater than 1%, and natural pregnancies miscarry at a rate greater than 10%. Assisted reproductive technologies should have a higher bar for safety, but by no means should they be required to produce 100% perfect results. Once they are known to improve over the other options, they should be available. We don't need to get to 1, but more like 0.99 (not speaking precisely).

So in this subarea of biology, we're not trying to go from 0 to 1, but more like from 0.01 to 0.99.

Biology's *scopes problem

Biology has had a *scopes problem[24]. Cells are tiny and numerous, and their contents are tinier and more numerous, so it's hard to know what all is going on between and inside cells.



(I don't know if this is how this meme works but I don't care and you can't stop me.)

This means that cell types, such as pluripotent stem cells or primordial germ cells or gametes or gametogonia, are not fully characterized; we can't say exactly or fully what it means for a cell to be that type of cell. Also, it's just expensive to measure large amounts of information about cells. Instead scientists use a few markers

[24]Adam Green. 'A Future History of Biomedical Progress', 1 August 2022. https://markovbio.github.io/biomedical-progress/.

associated with biotic processes to identify cells during experiments, such as the SYCP3 protein for meiosis, and then later take more costly comprehensive measurements to confirm results.

Measurement tech is developing quickly. E.g. DNA sequencing at scale has only been around for about 15 years, and single-cell RNA sequencing has spread in the past 10 years. Epigenomic sequencing has followed suit (e.g. bisulfite methylation sequencing)[25]. But these technologies are still far from ideal lenses; for example, the standard method for bisulfite sequencing requires applying the bisulfite treatment before amplification (because amplification by default produces unmethylated DNA), and it involves heating the DNA and therefore destroying much of it, so it has very poor coverage for single-cell sequencing. We coarsely point at types of cells by their physiological function, but we haven't yet carefully determined the full meaning of the categories.

Furthermore, many cells are difficult or infeasible to access. Many cells are embedded deep inside 3D tissues, maybe mixed in with an overwhelming majority of cells of some other type. Also, many tissues are either expensive, illegal, or unethical to access, most importantly tissue from humans such as fetal tissue or tissue from a living human's gonads. These cells and tissues are obviously much harder to characterize—we don't know what they look like naturally because we can only study a few examples (e.g. tissue from aborted fetuses or ovary tissue extracted from a woman undergoing cancer treatment for fertility preservation).

All of this means that terms like "secondary oocyte" or "spermatogonium" or "embryonic stem cell" are not 100% perfectly specified terms that point to a comprehensive catalog of known functional behaviors and internal states. Rather, they're phenomenological terms, as in "when transplanted into a conceptus, this kind of cell's descendants can contribute to any fetal tissues but not to the placenta, and maybe we have some noisy mixed RNA-seq data that somewhat characterizes some portion of the gene activity of this kind of cell". So a claim like "we created oocytes / oocyte-like cells" does not necessarily mean "we created cells that are fully competent to contribute to normal healthy embryonic development", and even if we did create fully competent cells we wouldn't be able to fully *tell* that we'd done so.

Methods to handle epigenomic correctness

There are two basic ways to handle epigenomic correctness:

- Somehow *make* your cell have the right epigenomic state. These are called "epigenomic correction" / "EC-making" methods.
- Let nature do the work of setting up the right epigenomic state, and do GV in a way that doesn't mess up nature's work. These are called "EC bypass" methods.

The following subsections describe the main types of methods to make cells that satisfy epigenomic correctness. There may be others I'm not aware of. Keep in mind that whether or not a method works for making cells that are *actually* epigenomically competent to reliably make a healthy human baby, there's a whole separate issue of *knowing* that it does so, and *demonstrating* that fact to scientists, regulators, and parents. That issue requires other research such as getting a clearer picture of epigenomically correct states, and validating EC-making methods in animal models to a high enough degree of reliability.

My beliefs about this area are still in flux, so take my claims with salt.

Takeaways:

- Full epigenomic correction in vitro—establishing full maternal imprinting or paternal imprinting—has not been achieved in humans. It also probably hasn't been achieved in mice, in the full sense of getting all the imprinting to look natural.
- Workable in vitro gametogenesis has been achieved in mice: We can make healthy mouse offspring from artificial gametes. But those methods won't translate very well to humans because they use tissue taken from mice, and they don't necessarily meet a high bar of epigenomic correctness.
- The robust way to accelerate the field:
  - Fund and do research that builds multi-omic single-cell atlases of human and primate reproduction. Only in the past 5 or so years has it even been possible to do so, because the technology was only recently developed[26].

[25] Baysoy, Alev, Zhiliang Bai, Rahul Satija, and Rong Fan. 'The Technological Landscape and Applications of Single-Cell Multi-Omics'. Nature Reviews Molecular Cell Biology 24, no. 10 (October 2023): 695–713. https://doi.org/10.1038/s41580-023-00615-w.

[26] Baysoy, Alev, Zhiliang Bai, Rahul Satija, and Rong Fan. 'The Technological Landscape and Applications of Single-Cell Multi-Omics'. Nature Reviews Molecular Cell Biology 24, no. 10 (October 2023): 695–713. https://doi.org/10.1038/s41580-

  - Fund and create a primate research center that is able to support tests of novel assisted reproductive technology.
- Instead of full EC-making, it may be easier to piggyback on natural gametogenesis. Chromosome selection on gamete DNA might be feasible, would bypass the EC problem, and would be a strong GV method. Iterated CRISPR editing SSCs followed by in vivo transplantation would be a medium-strength GV method.

In vitro gametogenesis (IVG)

Gametes are haploid cells (23 chromosomes, one of each index) that combine to form an embryo; male gametes are sperm (i.e. spermatozoa), female gametes are eggs (i.e. mature oocytes). Gametogenesis is the process whereby stem cells differentiate into gametes. In the natural lifecycle, the germline develops from early on in embryonic growth, culminating in gametogenesis in gonads (adult testicles, or fetal and adult ovaries).

In vitro gametogenesis (IVG) would reconstitute this process in the lab, making artificial sperm (in vitro spermatogenesis, IVS) or artificial eggs (in vitro oogenesis, IVO). IVG would allow us to artificially make sperm or eggs from stem cells that aren't just the natural germline stem cells that are native to humans.

Takeaways:

- Minimum viable IVO and IVS have both been achieved in mice, albeit using methods that wouldn't scale in humans because they use gonadal tissue extracted from organisms.
- Meiosis has been achieved, to some extent, in human male germline-cell-like cells. It hasn't been achieved in female germline-cell-like cells.
- Most research in human IVG so far uses gonadal tissue, so it wouldn't scale.
- Neither paternal EC-making nor maternal EC-making has been achieved in human cells in vitro.
- We need more sequencing data from natural gametogenesis.
  - E.g. scRNA-seq, scATAC-seq, scChIP-seq, and bisulfite seq atlases from human gonadal tissue. See e.g. [27] [28] [29] [30].
  - We need to know what natural gametes look like, so we know what the results of IVG should look like.
  - If we knew what natural gonadal cells—e.g. Sertoli cells, granulosa cells, thecal cells, and germline cells at various stages—normally look like in terms of gene expression, then we'd be able to coax iPSCs to behave likewise. Thus, we'd progress through germline development, and also make gonadal organoids able to support natural gametogenesis.
  - This might just work straightforwardly. IVO, and especially IVS, might or might not also need 3D culture methods. Testes in particular have 3D structure that supports sequential steps of development.
- Better methods for culturing gonadal tissue, over long periods and with high quality, would give us a way to cross the EC-making gap that could be accessible sooner than fully artificial end-to-end IVG.

Gametogenesis, and the research about it, is complex, and I am nowhere near to being an expert. The following subsections summarize some main points. For more reliable and complete information, see the reviews by Saitou and Miyauchi (2016)[31], Saitou and Hayashi (2021)[32], Tanaka and Watanabe (2023)[33],

023-00615-w.

[27] Guo, Jingtao, Edward J. Grow, Hana Mlcochova, Geoffrey J. Maher, Cecilia Lindskog, Xichen Nie, Yixuan Guo, et al. 'The Adult Human Testis Transcriptional Cell Atlas'. Cell Research 28, no. 12 (December 2018): 1141–57. https://doi.org/10.1038/s41422-018-0099-2.

[28] Sohni, Abhishek, Kun Tan, Hye-Won Song, Dana Burow, Dirk G. de Rooij, Louise Laurent, Tung-Chin Hsieh, et al. 'The Neonatal and Adult Human Testis Defined at the Single-Cell Level'. Cell Reports 26, no. 6 (5 February 2019): 1501-1517.e4. https://doi.org/10.1016/j.celrep.2019.01.045.

[29] Wang, Si, Yuxuan Zheng, Jingyi Li, Yang Yu, Weiqi Zhang, Moshi Song, Zunpeng Liu, et al. 'Single-Cell Transcriptomic Atlas of Primate Ovarian Aging'. Cell 180, no. 3 (6 February 2020): 585-600.e19. https://doi.org/10.1016/j.cell.2020.01.009.

[30] Jones, Andrea S. K., D. Ford Hannum, Jordan H. Machlin, Ansen Tan, Qianyi Ma, Nicole D. Ulrich, Yu-chi Shen, et al. 'Cellular Atlas of the Human Ovary Using Morphologically Guided Spatial Transcriptomics and Single-Cell Sequencing'. Science Advances 10, no. 14 (5 April 2024): eadm7506. https://doi.org/10.1126/sciadv.adm7506.

[31] Saitou, Mitinori, and Hidetaka Miyauchi. 'Gametogenesis from Pluripotent Stem Cells'. Cell Stem Cell 18, no. 6 (June 2016): 721–35. https://doi.org/10.1016/j.stem.2016.05.001.

[32] Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.

[33] Tanaka, Atsushi, and Seiji Watanabe. 'How to Improve the Clinical Outcome of Round Spermatid Injection (ROSI) into the Oocyte: Correction of Epigenetic Abnormalities'. Reproductive Medicine and Biology 22, no. 1 (9 February 2023): e12503. https://doi.org/10.1002/rmb2.12503.

Robinson et al. (2023)[34], and the other citations from this section. In what follows, I'll paint with a broad brush, glossing over very many potentially important details and distinctions.

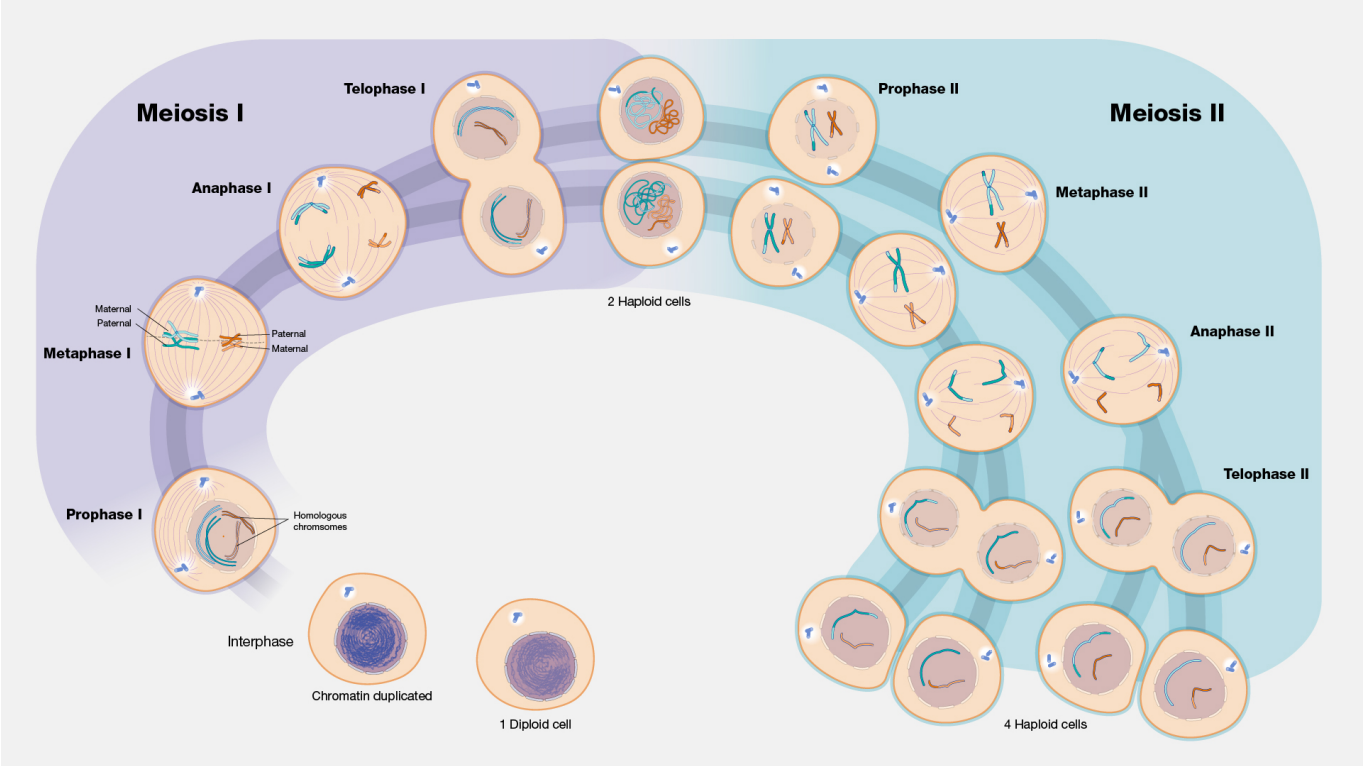
The basic elements of gametogenesis

The idea of IVG (in vitro gametogenesis) is to take some stem cells in a petri dish, and make them go through the cellular processes that happen in natural gametogenesis. You use chemicals (culture media, transcription factors, cytokines, gene editors) and surrounding cells to activate and support those cellular processes. Here are the three elements of gametogenesis, which are the changes that a stem cell should undergo to become a gamete:

1. make the needed epigenomic changes to the cell's DNA:
    a. general germ cell epigenomic reprogramming, first broadly wiping state and later silencing most of the genome;
    b. and sex-specific epigenetic imprinting at several dozen control sites;
2. perform meiosis, which makes a haploid cell with 23 chromosomes (one of each index) from a diploid cell with 46 chromosomes (two of each index) through recombination;
3. and make the cell develop, through sex-specific morphological and cytoplasmic changes.

The most important element for EC-making is 1., epigenomic correction. Unfortunately it's not very well understood, so I don't have a nice picture.

Meiosis looks like this:

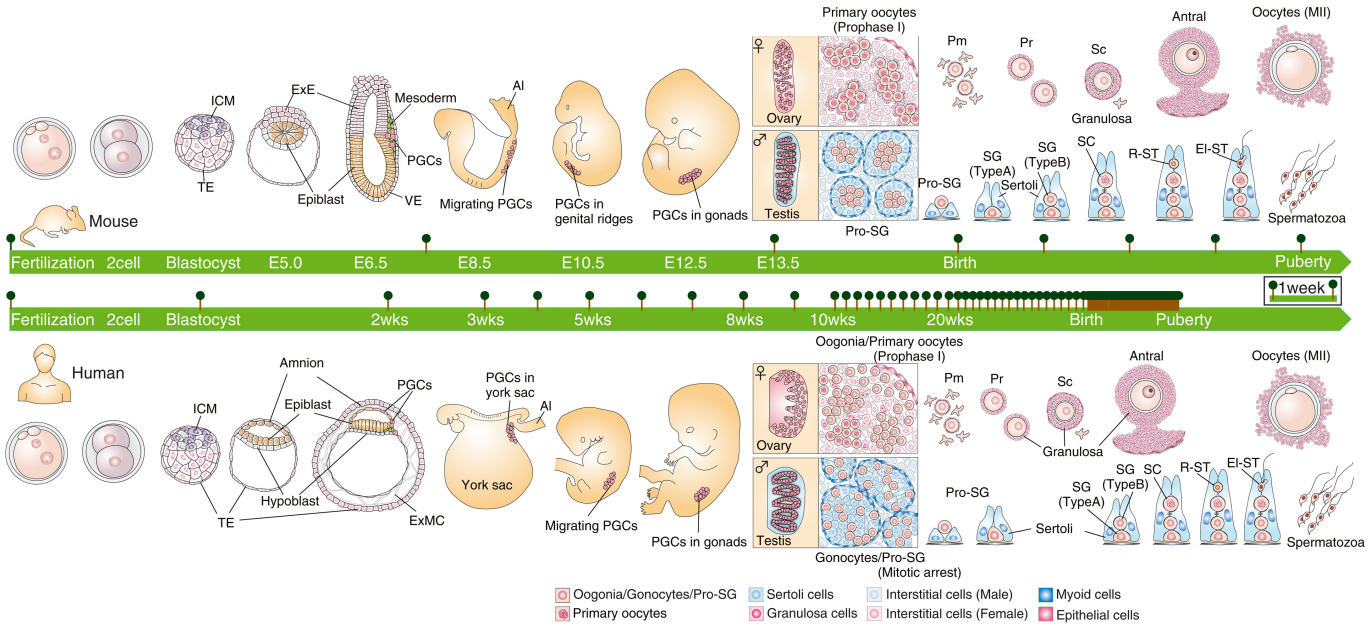


(Diagram from Gilchrist[35].)

Here's what mouse and human germline development, including gametogenesis, looks like:

[34]Robinson, Meghan, Sydney Sparanese, Luke Witherspoon, and Ryan Flannigan. 'Human in Vitro Spermatogenesis as a Regenerative Therapy — Where Do We Stand?' Nature Reviews Urology 20, no. 8 (August 2023): 461–79. https://doi.org/10.1038/s41585-023-00723-4.

[35]Gilchrist, Daniel A. 'Meiosis'. Accessed 2 March 2025. https://www.genome.gov/genetics-glossary/Meiosis.



(Figure 1 from Saitou and Hayashi (2021)[36].)

More zoomed in to oogenesis:

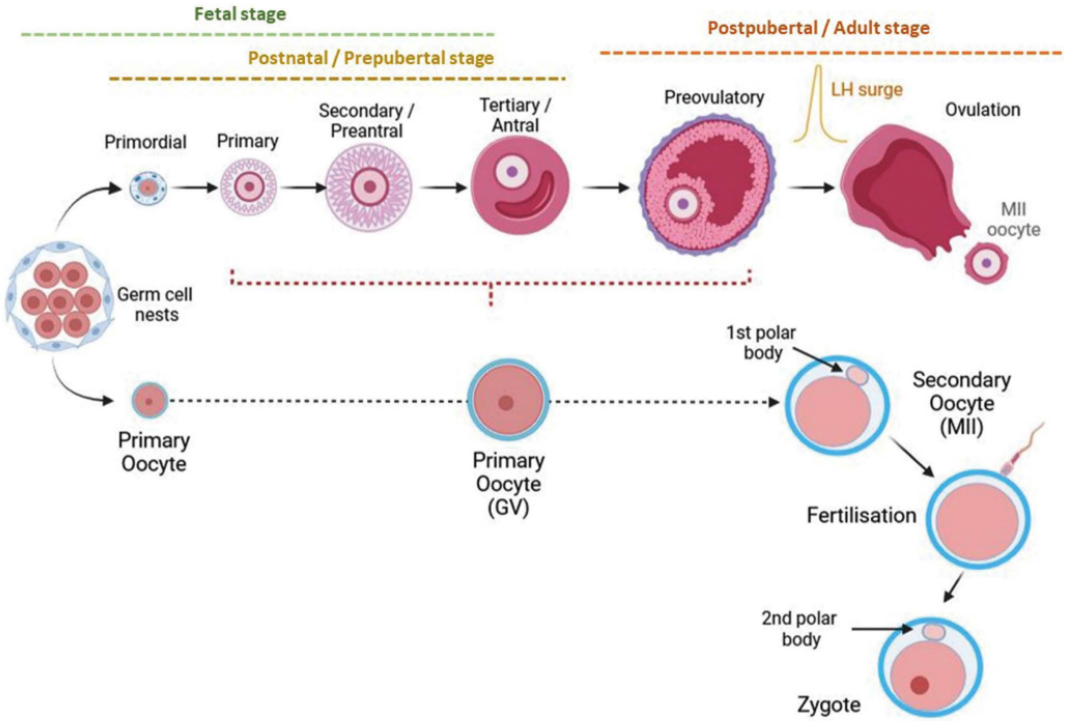


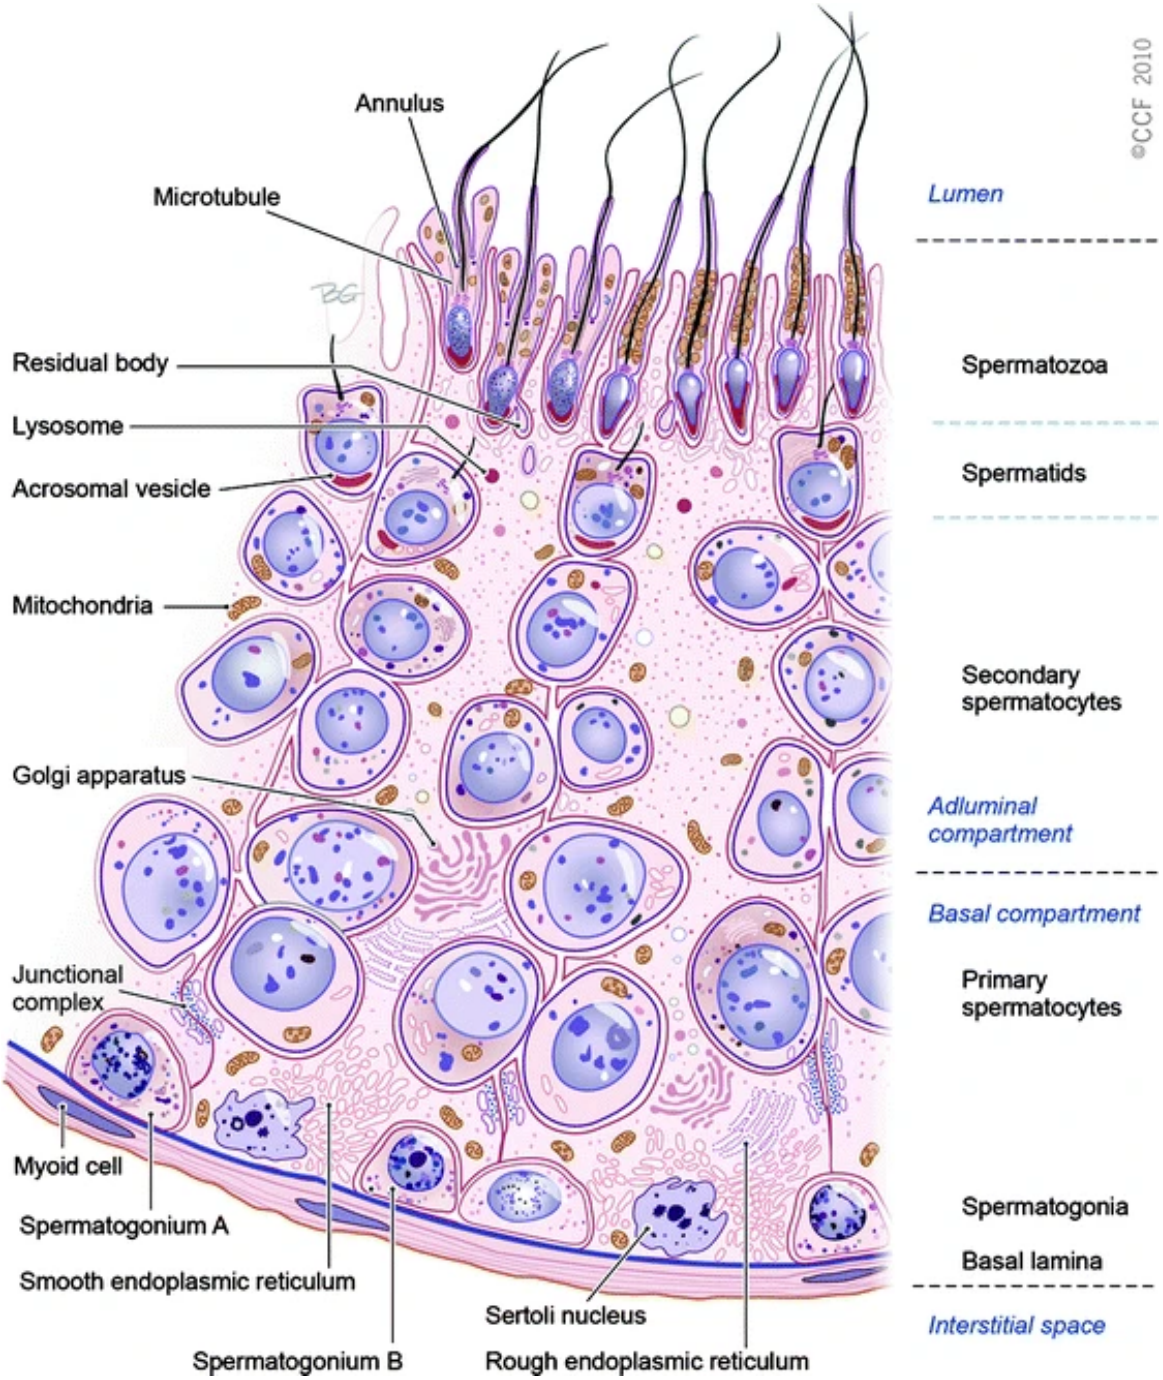
FIGURE 11.1 Schematic representation of the simultaneous events that encompass folliculogenesis and oogenesis. Early follicle formation takes place during fetal development with folliculogenesis progressing throughout female life, but final follicle growth and ovulation only occur after puberty, with the preovulatory LH surge. Diagram created using BioRender.com.

(Figure 11.1 from chapter 11 of Campell and Maalouf 2024[37].)

A depiction of spermatogenesis (the cells move from the bottom to the top):

[36]Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.

[37]Campbell, Alison, and Walid Maalouf. Mastering Clinical Embryology: Good Practice, Clinical Biology, Assisted Reproductive Technologies, and Advanced Laboratory Skills. CRC Press, 2024, https://www.google.com/books/edition/Mastering_Clinical_Embryology/xkb3EAAAQBAJ.



(Figure 2.3 from Sharma and Agarwal (2011)[38].)

IVG is a higher bar than epigenomic correction

A fully successful IVG method takes a stem cell, which could have been genomically vectored beforehand, and creates from it a competent gamete. You could take a somatic cell from an adult, apply your IVG method to make, say, an egg, and then fertilize that egg and make a healthy baby. This is a high bar, and has to fully complete the above three elements of gametogenesis: epigenomic reprogramming, meiosis, and cell development.

To qualify as an EC-making (epigenomic correction) method, an IVG protocol just has to make an *epigenomically* competent cell. The cell has to have nuclear DNA that has the right on-DNA sex-linked imprints and genome-wide reprogramming. Meiosis and cell development, while important, are not strictly necessary for EC-making.

The three basic elements of gametogenesis are necessary for fully competent oocytes. In natural oogenesis, oocytes grow quite large (>10x larger than progenitor oogonia) with extranuclear cytoplasm and accumulate a lot of cytoplasmic material—mitochondria and ribosomes, metabolic and regulatory proteins, and a large and diverse set of mRNAs. The size and contents of the cytoplasm kickstart and support the

[38]Sharma, Rakesh, and Ashok Agarwal. 'Spermatogenesis: An Overview'. In Sperm Chromatin: Biological and Clinical Applications in Male Infertility and Assisted Reproduction, edited by Armand Zini and Ashok Agarwal, 19–44. New York, NY: Springer, 2011. https://doi.org/10.1007/978-1-4419-6857-9_2.

development of the very early embryo while the zygote DNA is still largely silenced. For example, in very early embryonic growth, stored ribosomes produce proteins by translating stored mRNAs[39][40]. Meiosis is necessary so that the oocyte genome is haploid.

Only epigenomic correctness is strictly necessary for maternal EC-making. For many reproductive genomic vectoring purposes, it would minimally suffice to use donor eggs. The donor egg could come from a woman trying to have a child via reproductive GV, for herself, or it could come from a woman not otherwise involved in having the child. You take the nucleus (which contains the nuclear DNA of the donor) out of the donor egg, and you inject into the egg a nucleus containing DNA that's competent as maternal nuclear DNA—i.e., it is epigenomically correct. The donor egg provides the size and cytoplasmic material needed for the early embryo to grow. (In theory, one could even use a diploid cell with maternal imprinting, e.g. by using chromosome extraction to make a haploid maternally-EC cell from the diploid.)

In fully natural reproduction, all the elements of spermatogenesis are necessary. Natural spermatozoa (mature sperm) have, besides their DNA, several structures that are important for natural fertilization. A natural spermatozoon has, for example, mitochondria and a tail for swimming; an acrosome to break through the egg's outer barrier; and a centrosome, which helps organize the DNA for the zygote's first division.

But, like with the maternal germline, **only epigenomic correctness is necessary for paternal EC-making.** Intracytoplasmic sperm injection (ICSI), where you inject a sperm cell directly into the egg, seems to work ok. The injected cell doesn't need a tail (no swimming involved), and skips the acrosome reaction: the egg's outer barrier doesn't have to be broken down, since you're smuggling the sperm inside manually. The centrosome may be necessary, and the egg may have to be activated somehow (there are methods to do so, though they may be inefficient). Further, although oocytes contribute the lion's share of cytoplasmic material needed for early development, sperm may also contribute some important RNAs[41]. But sperm can if necessary be enucleated[42]. Since, unlike oocytes, natural healthy sperm are easy and inexpensive to obtain, there's no issue with using donor elements from sperm.

That said, even as a narrow EC-making method, the approach of IVG is to approximately recapitulate the natural process of oogenesis or spermatogenesis. Through this lens, the idea is that we are using nature's evolved gene regulatory mechanisms to make the maternal or paternal on-DNA epigenetic imprints, and the genome-wide reprogramming. So even if cell development and meiosis aren't strictly necessary for EC-making, **the straightforward approach to IVG as an EC-making method is to recapitulate all or almost all of natural gametogenesis**.

(In any case, a full IVO method that produces competent oocytes would be great. Besides being a breakthrough treatment for female infertility and for male-male couples, IVO would provide abundant and eventually inexpensive eggs. On its own, abundant eggs would perhaps double the effects of the existing GV method: simple embryo selection. Abundant eggs also makes any subsequent reproductive GV procedures less expensive and more powerful (more chances to try the procedure). Likewise, full IVS would inexpensively give fertility to infertile men and to female-female couples.)

Other general remarks about IVG

- IVG should be possible.
  - We know that the cell processes 1, 2, and 3, described above, are within a pluripotent stem cell's capabilities, because those processes operate in natural gametogenesis.
  - Activating these processes should be possible: as long as you replicate the natural signaling environment, support environment, and initial cell state well enough, the stem cell should behave just like it would in vivo. (There may also be gene regulatory shortcuts that jump straight to more developed cell states, skipping intermediate cell states.) If the process mimics natural

[39]Mtango, Namdori R., Santhi Potireddy, and Keith E. Latham. 'Chapter 7 Oocyte Quality and Maternal Control of Development'. In International Review of Cell and Molecular Biology, 268:223–90. Academic Press, 2008. https://doi.org/10.1016/S1937-6448(08)00807-1.

[40]Pors, S. E., D. Nikiforov, J. Cadenas, Z. Ghezelayagh, Y. Wakimoto, L. A. Z. Jara, J. Cheng, et al. 'Oocyte Diameter Predicts the Maturation Rate of Human Immature Oocytes Collected Ex Vivo'. Journal of Assisted Reproduction and Genetics 39, no. 10 (October 2022): 2209–14. https://doi.org/10.1007/s10815-022-02602-0.

[41]Champroux, Alexandre, Julie Cocquet, Joëlle Henry-Berger, Joël R. Drevet, and Ayhan Kocer. 'A Decade of Exploring the Mammalian Sperm Epigenome: Paternal Epigenetic and Transgenerational Inheritance'. Frontiers in Cell and Developmental Biology 6 (15 May 2018). https://doi.org/10.3389/fcell.2018.00050.

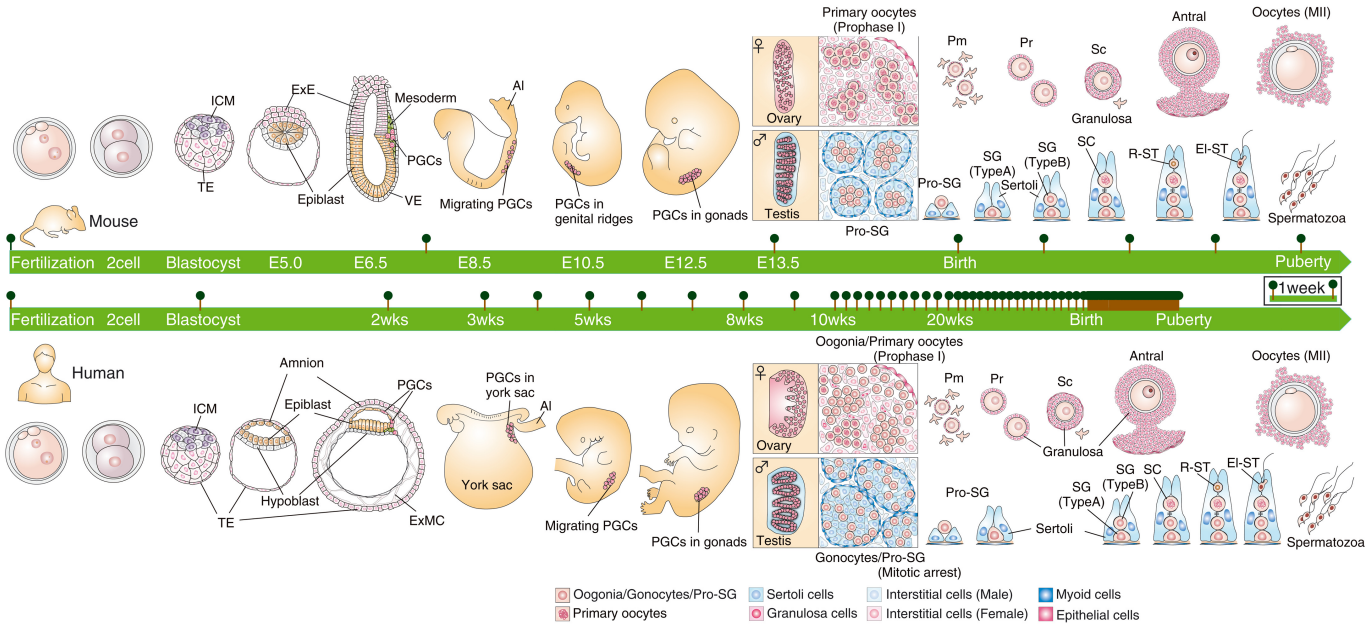
[42]Yoon, Sook-Young, and Rafael A. Fissore. 'Release of Phospholipase C Zand [Ca2+]i Oscillation-Inducing Activity during Mammalian Fertilization', 1 November 2007. https://doi.org/10.1530/REP-07-0259.

gametogenesis well enough, the resulting haploid cell will be able to play the role of a natural gamete.

- Academic IVG research generally says that it is motivated by understanding and potentially improving reproductive health, which is a different motivation from reproductive genomic vectoring.
  - Presumably researchers would be quite happy if they could actually generate human gametes competent to make babies, in order to have a general solution to infertility. But a lot of research is aimed at having any model at all, in vitro, of various steps in gametogenesis. Even a slow, expensive, inefficient model can be useful for studying gametogenesis, e.g. to understand the cellular mechanisms involved and maybe to treat specific infertility problems with simple chemical treatments. While this is an admirable goal, and has been until recently the sole source of progress toward eventually having a more practical system for IVG, it's a somewhat different orientation to research.
  - **The goal relevant to germline engineering is more like: How can we recapitulate, in vitro, the epigenomic reprogramming that occurs in natural human gametogenesis—or mimic the final results—to such a high standard that we can actually use the resulting DNA to make babies?** Furthermore, how can we do so reliably and efficiently at scale? Separately, how can we perform meiosis, as is done in natural gametogenesis, but maybe by a different method? How can we accelerate these processes so they don't take many months? How can we make all this happen ASAP, rather than eventually?
  - **These questions induce different research priorities.** For example, IVG research aimed at solving fertility often assumes that, in the context of applications, naturally epigenomically corrected DNA is not available. But for reproductive GV, that assumption doesn't have to be made, so we can consider using naturally EC DNA, or piggybacking off of natural EC-making, or using a prospective parent's own healthy gonadal tissue in vitro.
- There's no clear line where a method qualifies as in vitro gametogenesis.
  - Besides performance (efficiency, cost), different methods have different inputs:
    - some kind of stem cell to be turned into oocytes (iPSCs, ESCs, PGCs, oogonia, oocytes, prospermatogonia, SSCs, spermatocytes);
    - some kind of gonad-like cells (tissue extract, organoids, reconstituted tissue, xeno gonads);
    - and some culture media.
  - Different methods also have different outputs: spermatogonia or spermatocytes or spermatids, or immature oocytes or mature oocytes, of varying qualities.
  - Because of the *scopes problem in biology, the meaning and measurement of germline cell quality is fuzzy. The highest standard would be to take a somatic cell, and then with no additional cells taken from any organism, produce a cell that is totaly indistinguishable from a natural mature oocyte or a natural spermatozoon.
- It is possible to get an IVG-like method that creates cells that are haploid, but that have only some of the needed imprinting.
  - So, even if we get a partially working IVG method, some understanding of epigenomic correctness is still necessary. In order to avoid an epigenomic near-miss, we'd need to know when our cell is in a good enough state to be a gamete or zygote. We can't just trust that, since something like gametogenesis happened, the result is a fully competent gamete.
  - In any case, a method for inducing meiosis, regardless of epigenetics, would be useful for iterated meiotic selection or recombinant chromosome selection.

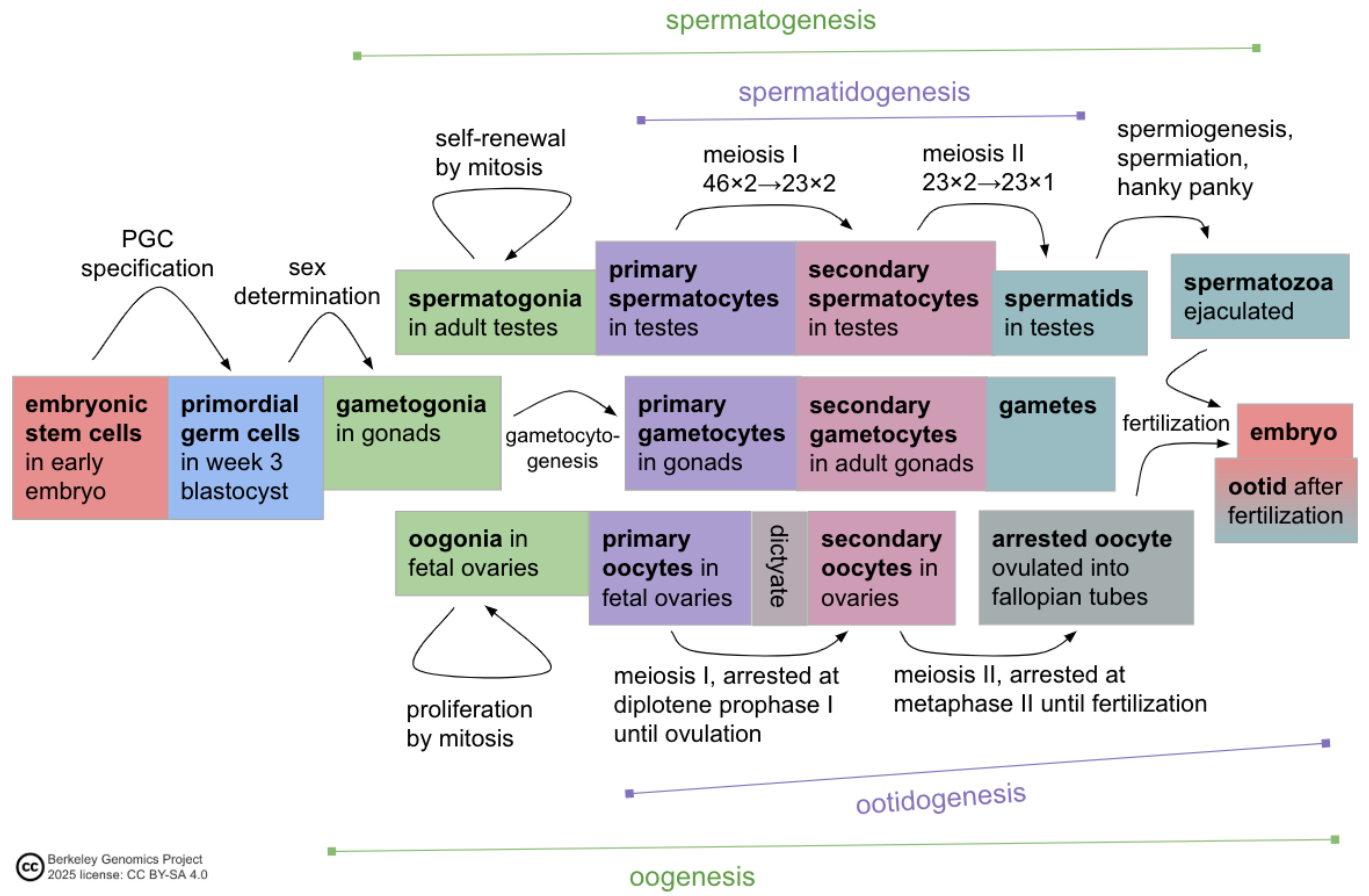
The general state of IVG research

We have to rewind a bit before talking about gametogenesis proper. Here's the beautiful illustration we saw earlier, showing the process of natural mammalian germline development:



(Figure 1 from Saitou and Hayashi (2021)[43].)

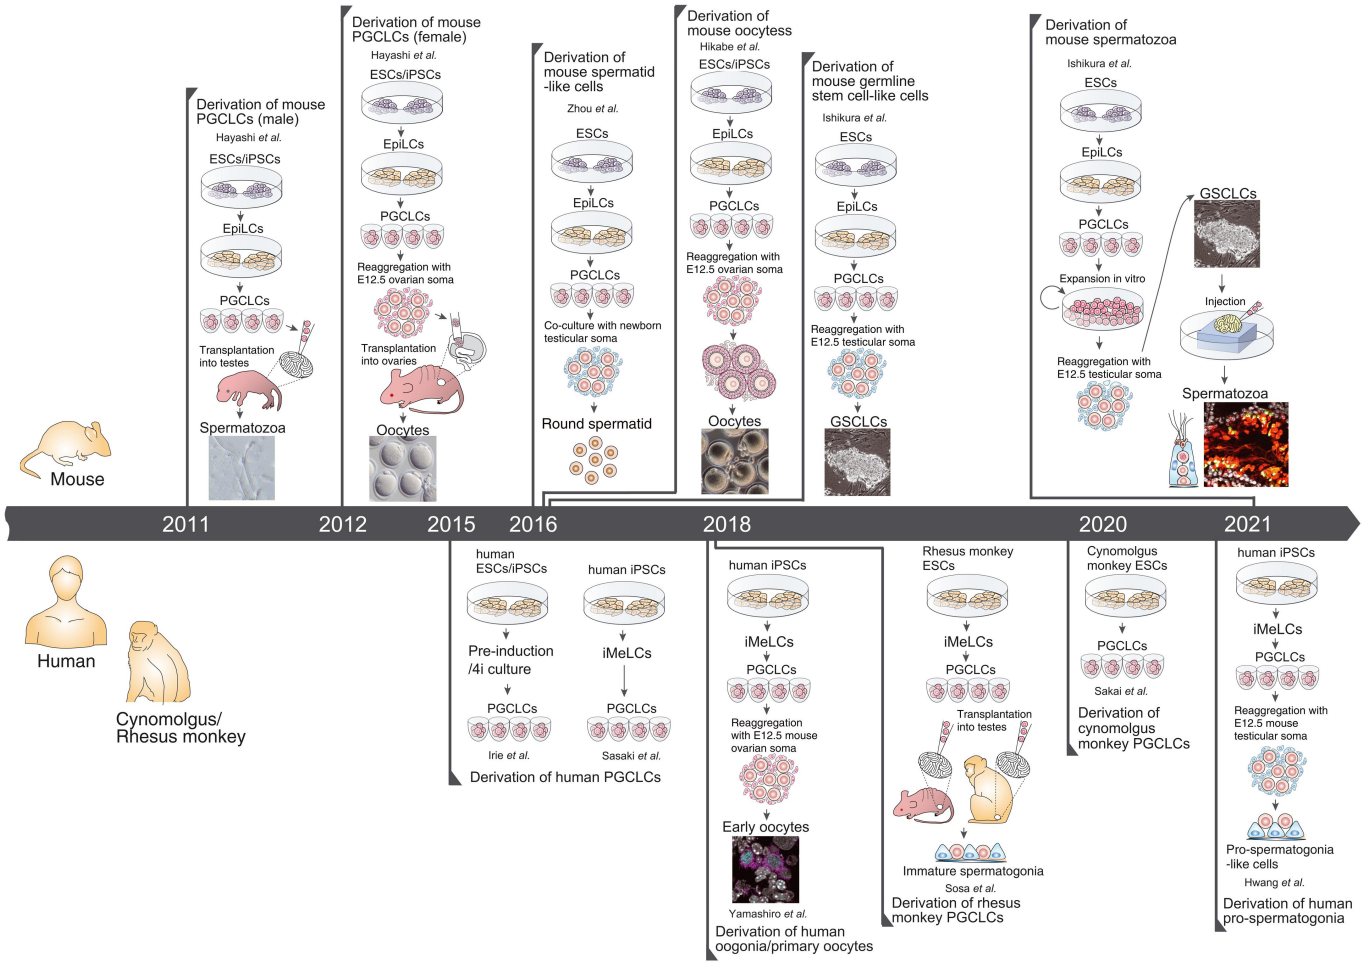
Here's a more schematic diagram, with terminology:



Gametogenesis (oogenesis, spermatogenesis) is shown between the faint green brackets.

Here's a summary of the state of IVG research from 2021:

[43]Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.



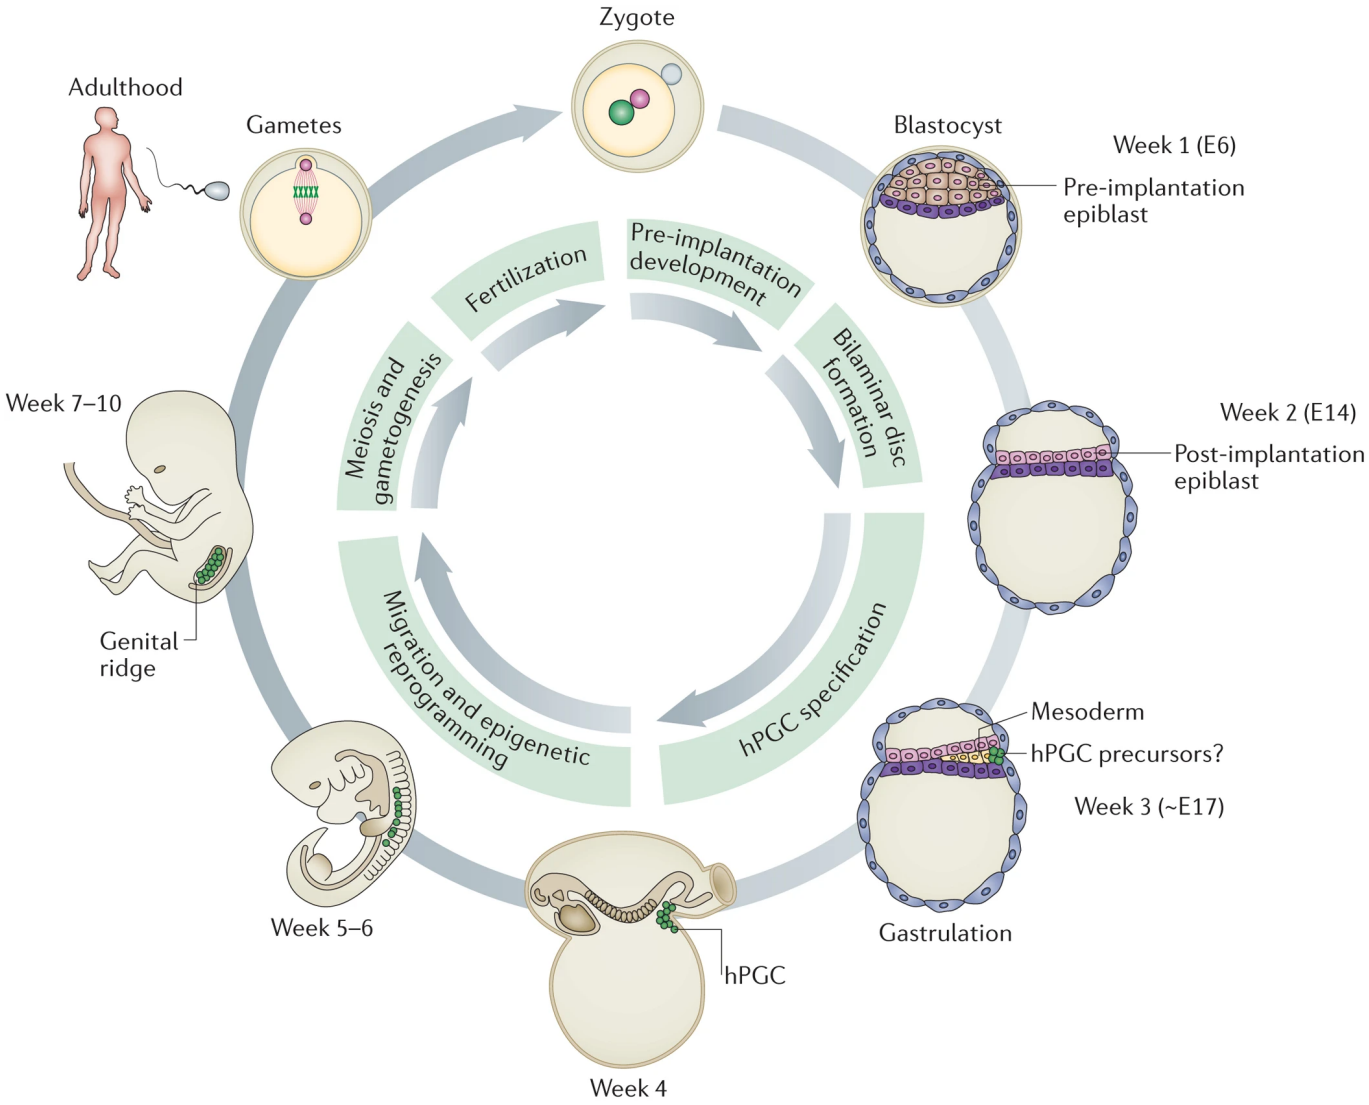
(Figure 4 from Saitou and Hayashi (2021)[44].)

Most of the progress has relied on culturing stem cells with some sort of special stem cells for support, such as embryonic stem cells or specifically gonadal stem cells, taken from mice or monkeys or humans. Chemical additives (cytokines, etc.) are important, but the support cells are as yet indispensable for the hardest steps of gametogenesis. This suggests that the likely way forward for human IVG will involve ovarian follicle culture, so that area ought to be afforded additional talented researchers; see the "Key challenges" section of Saitou and Hayashi (2021)[45].

Gametogonia are the stem cells that proliferate themselves by mitotically dividing, and that can also differentiate into gametes. Before gametogonia specialize to be female (oogonia) or male (spermatogonia), they have a shared developmental history as primordial germ cells (PGCs). PGCs form in the early embryo, then migrate to the developing gonads, and then differentiate into gametogonia. An illustration of the whole human germline cycle:

[44]Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.

[45]Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.



(Figure 1 from Tang et al., 2016[46].)

PGCs (primordial germ cells) have a distinct epigenomic state from other embryonic stem cells (ESCs). During germline specification, the ESCs that become PGCs have their epigenetic markers almost completely wiped, including sex-linked imprints, whereas other ESCs mostly keep their imprints as they differentiate to form other fetal tissues. Can we recapitulate this in vitro? Human PGC-like cells spontaneously form, with some low frequency, just from an in vitro culture of human ESCs[47].

We could use ESCs. But if possible, to avoid having to take cells from an embryo, it's better to use iPSCs. Induced pluripotent stem cells (iPSCs) come from somatic (non-stem) cells, taken from an organism's body. These somatic cells are already differentiated, and can't easily be converted into other cells, and generally can't even mitotically divide. Takahashi and Yamanaka (2006) discovered how to use a cocktail of four transcription factors (now called Yamanaka factors) to de-differentiate somatic cells into pluripotent stem cells, which can divide and can potentially differentiate into many kinds of cells[48].

Irie et al. (2015) (using their 4i medium with four inhibitory cytokines), as well as Sasaki et al. (2015), found fairly efficient methods to make human PGC-like cells (PGCLCs) from human ESCs or from human iPSCs

[46]Tang, Walfred W. C., Toshihiro Kobayashi, Naoko Irie, Sabine Dietmann, and M. Azim Surani. 'Specification and Epigenetic Programming of the Human Germ Line'. Nature Reviews Genetics 17, no. 10 (October 2016): 585–600. https://doi.org/10.1038/nrg.2016.88.

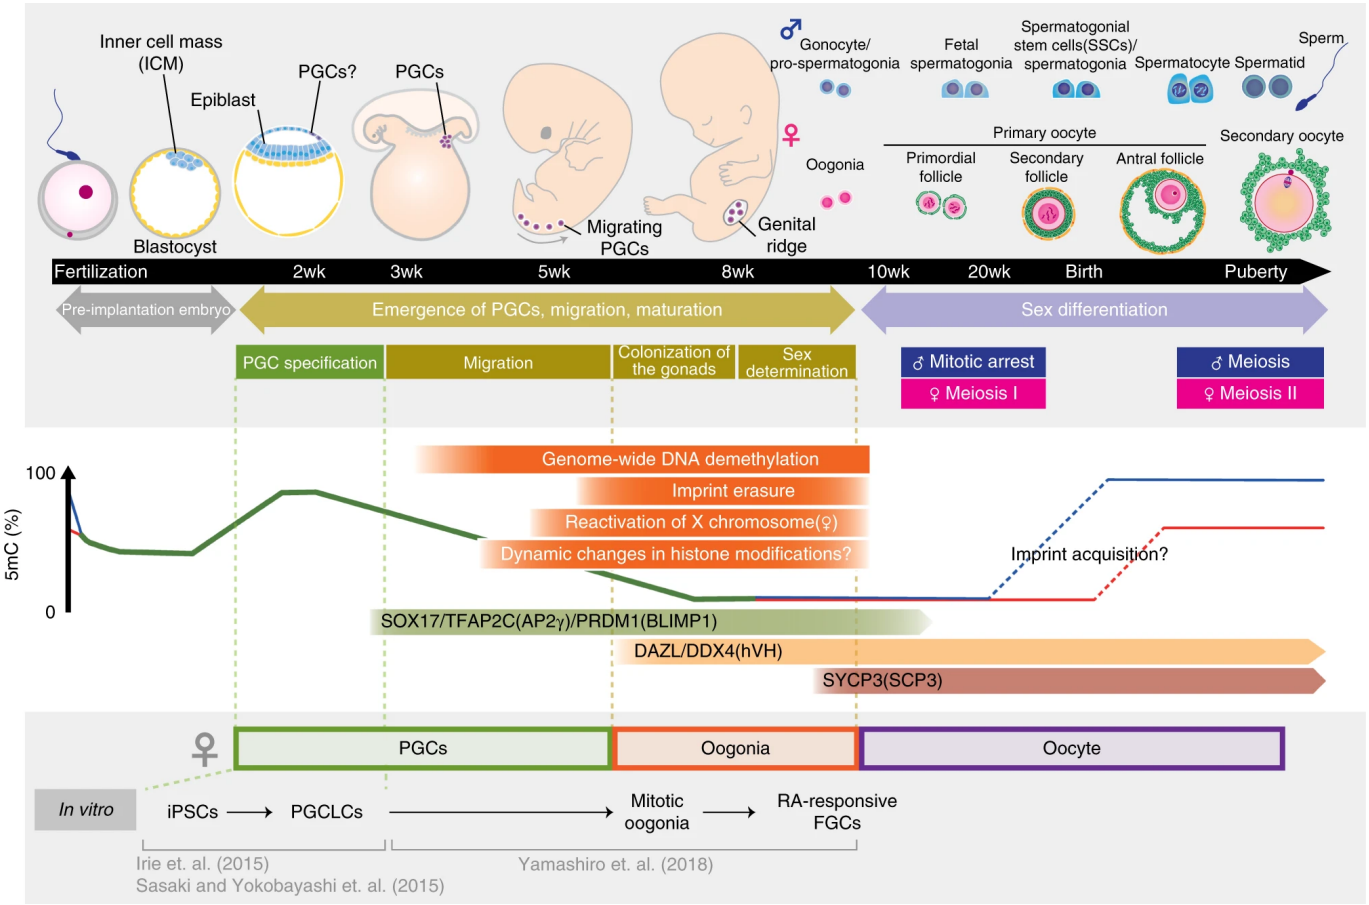
[47]Gkountela, Sofia, Ziwei Li, John J. Vincent, Kelvin X. Zhang, Angela Chen, Matteo Pellegrini, and Amander T. Clark. 'The Ontogeny of cKIT+ Human Primordial Germ Cells Proves to Be a Resource for Human Germ Line Reprogramming, Imprint Erasure and in Vitro Differentiation'. Nature Cell Biology 15, no. 1 (January 2013): 113–22. https://doi.org/10.1038/ncb2638.

[48]Takahashi, Kazutoshi, and Shinya Yamanaka. 'Induction of Pluripotent Stem Cells from Mouse Embryonic and Adult Fibroblast Cultures by Defined Factors'. Cell 126, no. 4 (25 August 2006): 663–76. https://doi.org/10.1016/j.cell.2006.07.024.

[49][50]. (See also Panula et al. (2011)[51].) Together, iPSC induction and iPGCLC induction wipe much of the epigenetic marks from the cell's DNA. This partially mimics the effect of natural PGC specification, and puts the cell in a very undifferentiated state, ready to be reprogrammed into gametes.

Once we're at the PGC (or PGC-like cell) stage, the next step is sex specification, which turns PGCs into gametogonia (oogonia, spermatogonia). The oogonia and spermatogonia will differentiate into primary gametocytes (respectively oocytes and spermatocytes). It's around this time that sex-specific imprinting is established in these germline cells. Finally, the gametocytes undergo meiosis, producing (eventually) gametes. Since these steps are sex-specific, they'll be addressed in their respective following subsections on IVO and IVS.

Here's an illustration of gametogenesis, including epigenomic reprogramming:



(Figure 1 from Yamashiro et al. 2020[52].)

Of special importance for our EC-making purposes is the developmental periods when epigenetic imprinting is established (middle right). The dynamics of imprinting in male and female germlines aren't fully understood, but broadly the picture seems to be:

- Paternal imprinting is established starting in prospermatogonia as they differentiate to spermatogonial stem cells, and is at least largely complete by the time the primary spermatocyte, in adulthood, is partway into meiosis I (pachytene substage of prophase I)[53][54].

[49]Irie, Naoko, Leehee Weinberger, Walfred W.C. Tang, Toshihiro Kobayashi, Sergey Viukov, Yair S. Manor, Sabine Dietmann, Jacob H. Hanna, and M. Azim Surani. 'SOX17 Is a Critical Specifier of Human Primordial Germ Cell Fate'. Cell 160, no. 1–2 (15 January 2015): 253–68. https://doi.org/10.1016/j.cell.2014.12.013.

[50]Sasaki, Kotaro, Shihori Yokobayashi, Tomonori Nakamura, Ikuhiro Okamoto, Yukihiro Yabuta, Kazuki Kurimoto, Hiroshi Ohta, et al. 'Robust In Vitro Induction of Human Germ Cell Fate from Pluripotent Stem Cells'. Cell Stem Cell 17, no. 2 (6 August 2015): 178–94. https://doi.org/10.1016/j.stem.2015.06.014.

[51]Panula, Sarita, Jose V. Medrano, Kehkooi Kee, Rosita Bergström, Ha Nam Nguyen, Blake Byers, Kitchener D. Wilson, et al. 'Human Germ Cell Differentiation from Fetal- and Adult-Derived Induced Pluripotent Stem Cells'. Human Molecular Genetics 20, no. 4 (15 February 2011): 752–62. https://doi.org/10.1093/hmg/ddq520.

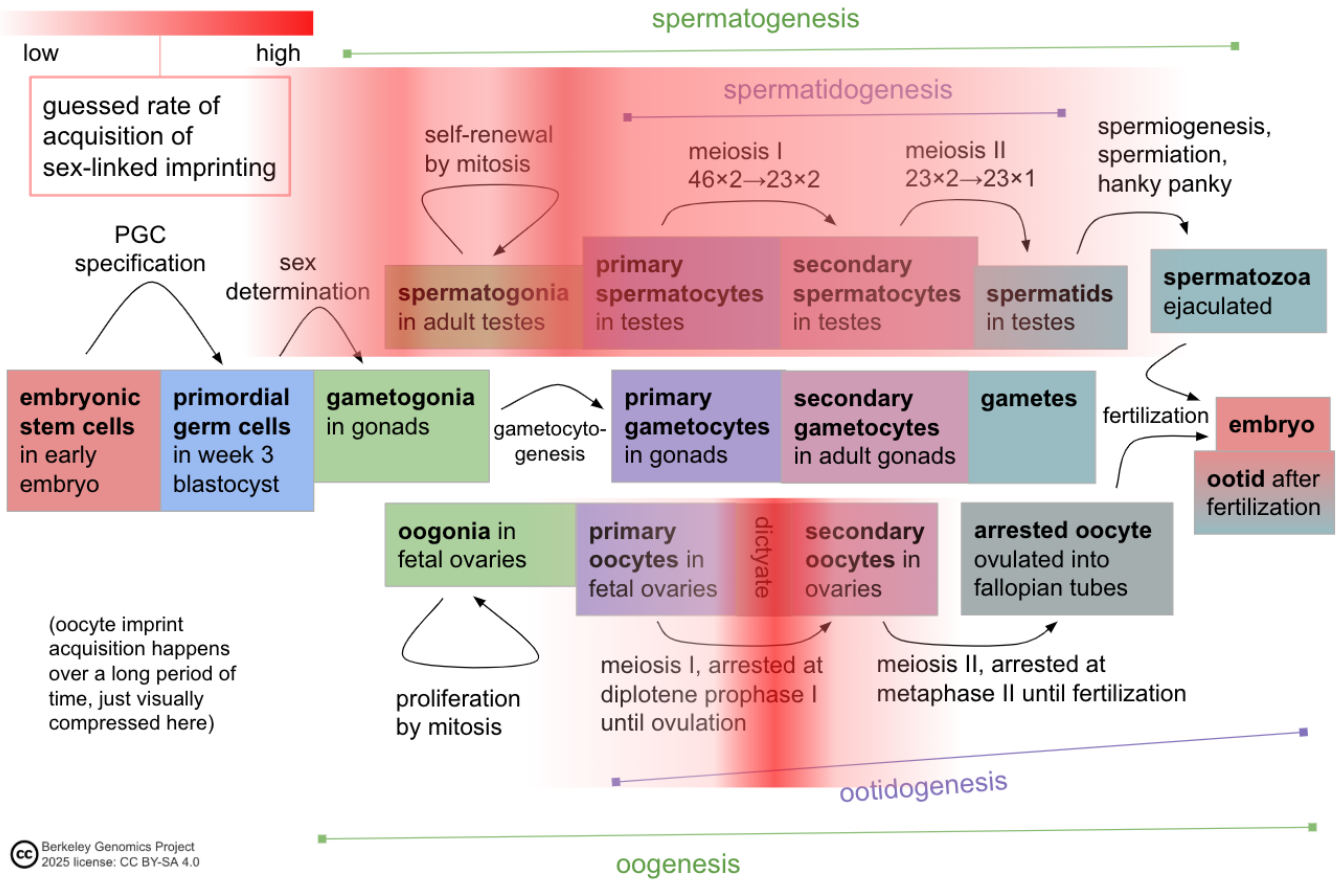
[52]Yamashiro, Chika, Kotaro Sasaki, Shihori Yokobayashi, Yoji Kojima, and Mitinori Saitou. 'Generation of Human Oogonia from Induced Pluripotent Stem Cells in Culture'. Nature Protocols 15, no. 4 (April 2020): 1560–83. https://doi.org/10.1038/s41596-020-0297-5.

[53]Trasler, Jacquetta M. 'Epigenetics in Spermatogenesis'. Molecular and Cellular Endocrinology, Proceedings of the 15th Workshop on Molecular and Cellular Endocrinology of the Testis, 306, no. 1 (10 July 2009): 33–36. https://doi.org/10.1016/j.mce.2008.12.018.

[54]Kelsey, Gavin, and Robert Feil. 'New Insights into Establishment and Maintenance of DNA Methylation Imprints in

- Maternal imprinting is established postnatally, in the slow prepubertal growth of the dictyate oocyte[55].

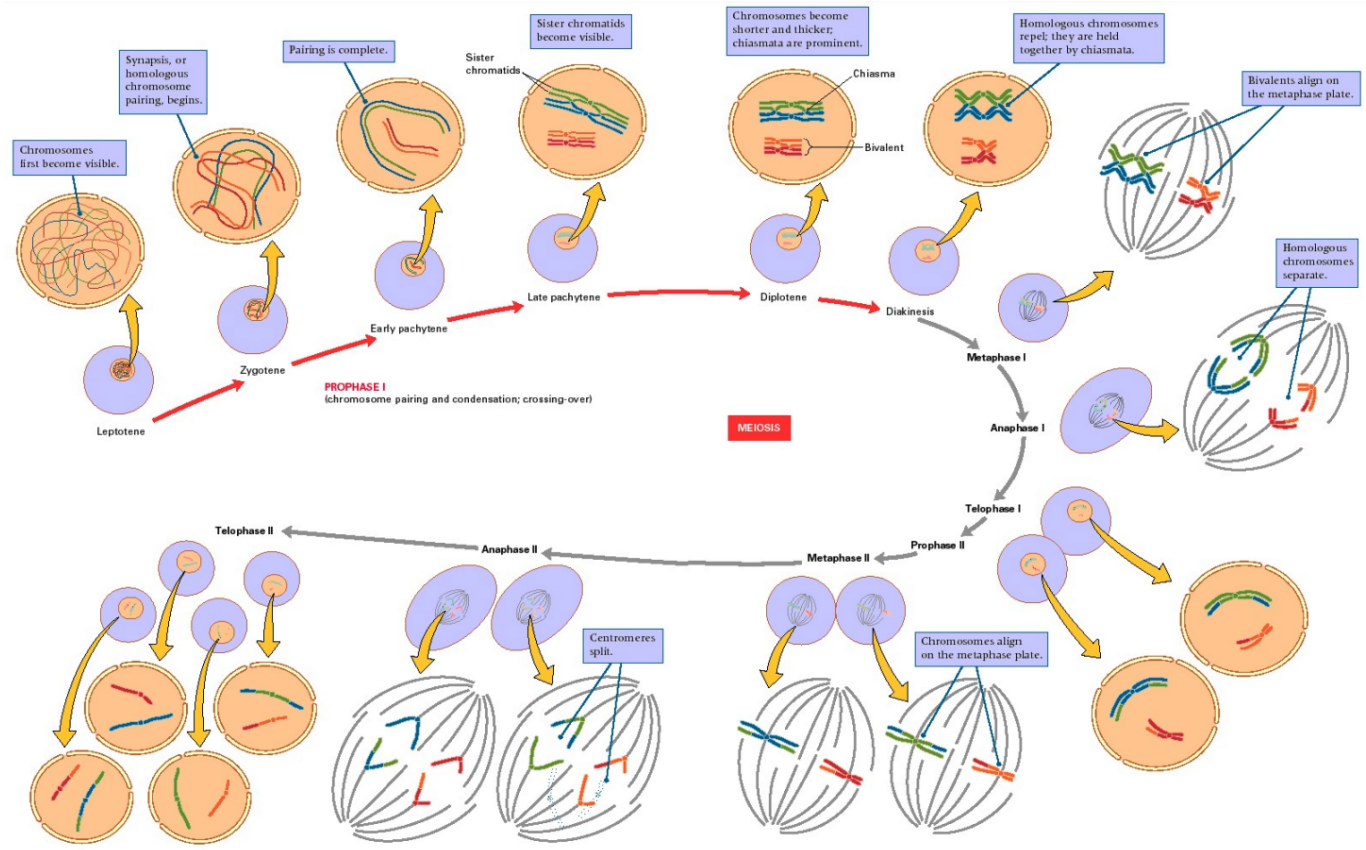
A schematic diagram showing when imprinting is believed to be established:



Both IVO and IVS will have to make the cell undergo meiosis. In meiosis, the diploid cell with 46 chromosomes copies its DNA (46 → 92), performs chromosome recombination, and then divides twice (92 → 46 → 23) to produce four haploid daughter cells with 23 chromosomes each. Meiosis is a complicated process that takes at least a few weeks in human cells (generally a few months in natural gametogenesis, depending how you count):

Mammals'. Philosophical Transactions of the Royal Society B: Biological Sciences 368, no. 1609 (5 January 2013): 20110336. https://doi.org/10.1098/rstb.2011.0336.

[55]Trasler, Jacquetta M. 'Gamete Imprinting: Setting Epigenetic Patterns for the next Generation'. Reproduction, Fertility and Development 18, no. 2 (2006): 63. https://www.publish.csiro.au/rd/Fulltext/RD05118.



(I'm unsure where this diagram is from originally; there are several versions floating around. Keser (2016)[56] cites it to Usserly (1998)[57].)

Ivy Natal is a commercial company aiming to achieve meiosis first in animals for agricultural purposes.

In fact, meiosis may happen spontaneously in cultured human stem cells; several projects have claimed to achieve meiosis in human cells, with only a few or no genetic tweaks[58] (see [59] for discussion). However, the efficiency of those methods is quite low, producing <2% haploid cells, and those studies might not replicate straightforwardly or at all. Further, it's unclear what kind of cells they made; some of them may have been dead or dying cells with reduced DNA content, rather than actual haploid cells produced by meiosis. Smela et al. (2024) demonstrate directly inducing human iPSCs to quickly (two weeks) get partway through prophase I of meiosis (maybe 0.1% up to the pachytene substage)[60]. Their method induced >15% of the starting stem cells to enter meiosis:

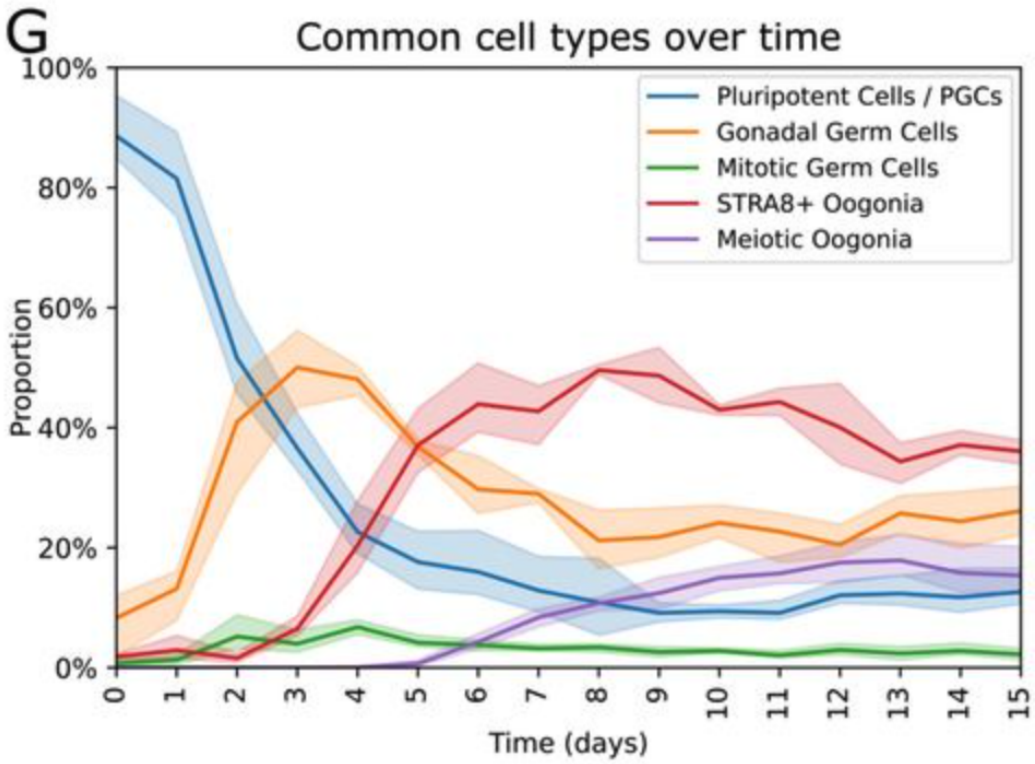
[56]Keser, Vanessa. 'Lokalizacija proteina kontrolne točke diobenog vretena u mišjim spermatocitama'. Info:eu-repo/semantics/bachelorThesis, University of Zagreb. Faculty of Food Technology and Biotechnology. Department of Biochemical Engineering. Laboratory for Biology and Microbial Genetics, 2016. https://urn.nsk.hr/urn:nbn:hr:159:787761.

[57]Usserly, D. (1998): Gene Exchange, Meiosis & Eukaryotic Life Cycles. In Biology 101: Life on Planet Earth, http://www.cbs.dtu.dk/courses/genomics_course/roanoke/bio101ch12.htm.

[58]Eguizabal, C., N. Montserrat, R. Vassena, M. Barragan, E. Garreta, L. Garcia-Quevedo, F. Vidal, A. Giorgetti, A. Veiga, and J. C. Izpisua Belmonte. 'Complete Meiosis from Human Induced Pluripotent Stem Cells'. STEM CELLS 29, no. 8 (2011): 1186–95. https://doi.org/10.1002/stem.672.

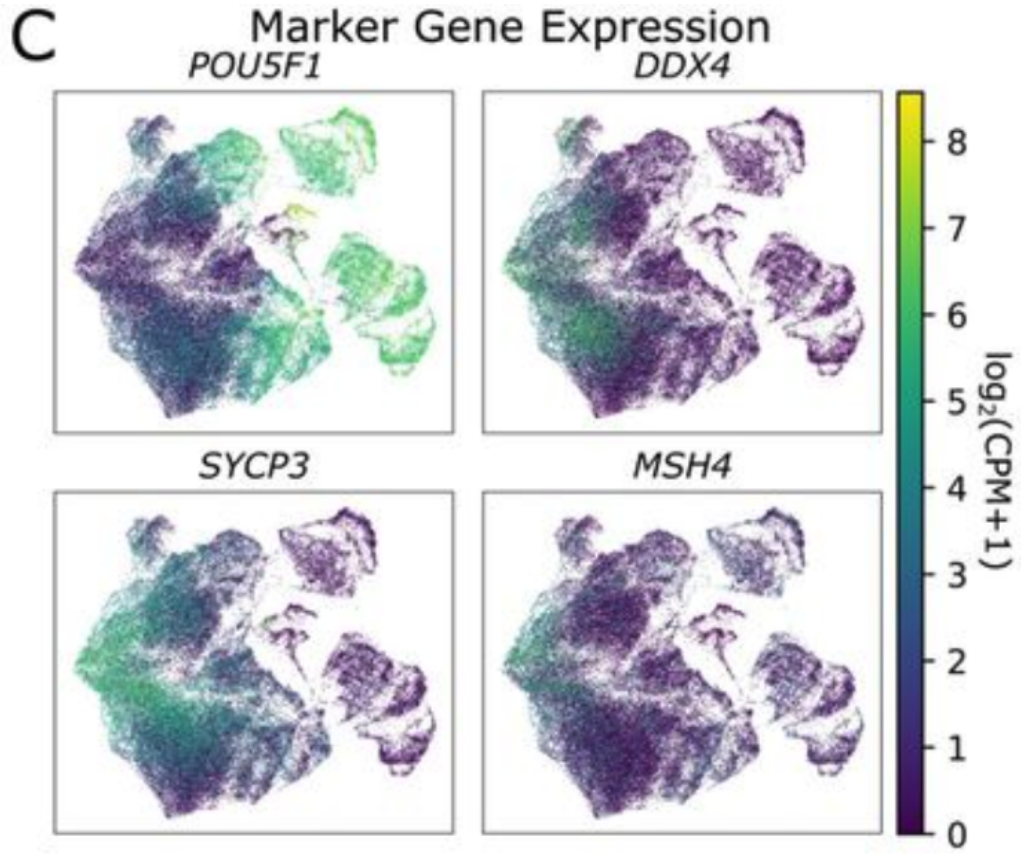
[59]Saitou, Mitinori, and Hidetaka Miyauchi. 'Gametogenesis from Pluripotent Stem Cells'. Cell Stem Cell 18, no. 6 (June 2016): 721–35. https://doi.org/10.1016/j.stem.2016.05.001.

[60]Pierson Smela, Merrick, Jessica Adams, Carl Ma, Laura Breimann, Ursula Widocki, Toshi Shioda, and George M. Church. 'Induction of Meiosis from Human Pluripotent Stem Cells'. bioRxiv, 31 May 2024, 2024.05.31.596483. https://doi.org/10.1101/2024.05.31.596483.



(Figure 5G from Smela et al. (2024).)

A small fraction of the cells expressed MSH4 (the greenish ones in the lower right), a protein that's involved specifically in the pachytene substage of prophase I, suggesting some cells reached that point:

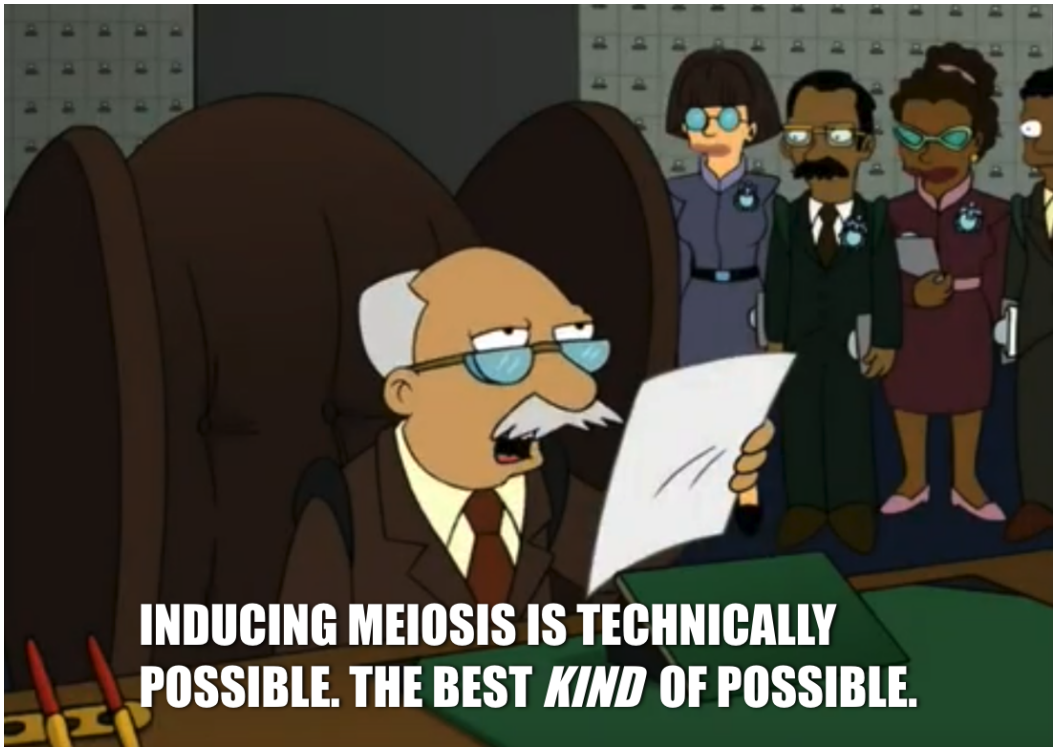


(Figure 5C from Smela et al. (2024).)

Further, by culturing fetal testis tissue, Yuan et al. (2020) produced haploid cells and more thoroughly verified their haploidy[61]. In principle, this provides a method for meiosis in vitro, maybe combined with

[61] Yuan, Yan, Laihua Li, Qing Cheng, Feiyang Diao, Qiao Zeng, Xiaoyu Yang, Yibo Wu, et al. 'In Vitro Testicular Organogenesis from Human Fetal Gonads Produces Fertilization-Competent Spermatids'. Cell Research 30, no. 3 (March 2020): 244–55. https://doi.org/10.1038/s41422-020-0283-z.

Hwang et al.'s (2020) method for inducing human prospermatogonia-like cells in vitro[62]. This would be an inefficient and non-scalable method (requires human tissue) for inducing meiosis, but shows that it's *technically* possible in some sense.



As will be mentioned below, both IVO and IVS have already been achieved in mice, in some fashion. But mouse reproductive epigenetics is different from human reproductive epigenetics, so results from mice don't necessarily translate to humans (or other mammals)[63]. Part of the reason that mice are a model organism for reproduction is that certain lab mouse cell lines are uniquely easy to coax into certain cellular processes. Also, there hasn't yet been an extensive study of the health of mice grown from eggs and sperm derived via IVO and IVS. Such a study would greatly increase the flow of valuable feedback about which epigenetic imprints are and are not important for completely healthy development, and which of those imprints are and are not corrected by IVG methods.

In vitro oogenesis (IVO)

In vitro oogenesis (IVO) means taking human stem cells and inducing them to differentiate into functionally normal oocytes ready to be fertilized with a sperm and thus contribute to a healthy baby.

Takeaways:

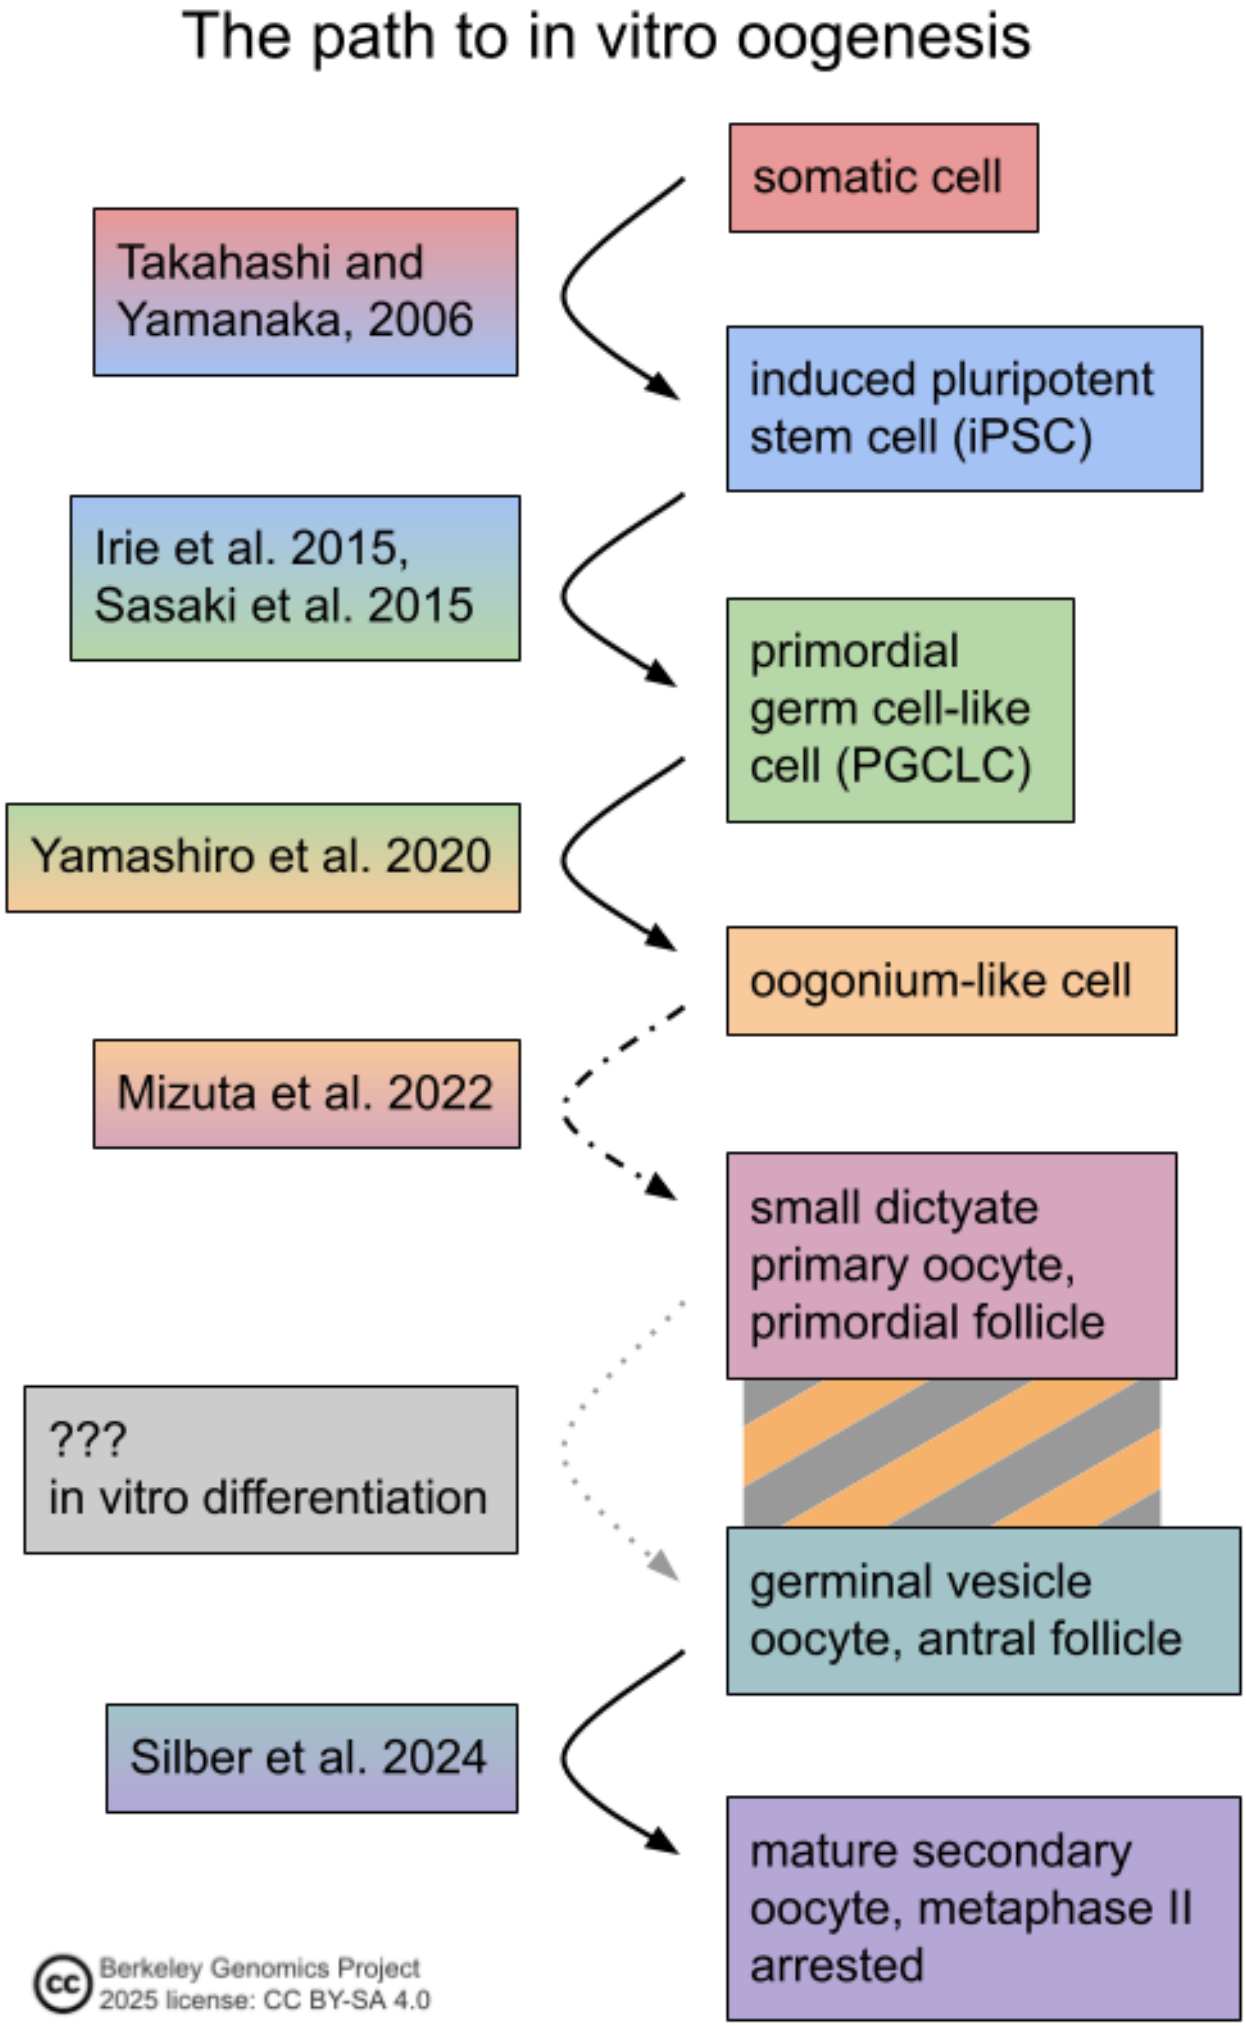
- IVO research has been making progress on oogenesis in the initial stages (up to primary oocytes) and final stages (final oocyte maturation). But there is a core uncrossed gap: maturing a small primary oocyte and primordial follicle from the perinatal stage to the adult stage, where it should be a larger germinal vesicle oocyte in a preantral / antral follicle. This step is crucial because—besides making the oocyte competent on its own to support the growth of the zygote—this step is where maternal epigenetic imprinting is established.
- The most successful methods for making primary oocytes in vitro use human fetal tissue, and therefore are not scalable. There is some progress on methods that instead use only iPSCs, culture media, and targeted gene overexpression.
- More data (scRNA, scATAC, scChIP, bisulfite seq) from natural oogenesis would tell us more about how to control gene expression to activate oogenesis.
- Better methods for culturing ovarian tissue or ovarian organoids, especially fetal and postnatal ovary-like tissue, might help us cross the maternal EC-making gap.

[62]Hwang, Young Sun, Shinnosuke Suzuki, Yasunari Seita, Jumpei Ito, Yuka Sakata, Hirofumi Aso, Kei Sato, Brian P. Hermann, and Kotaro Sasaki. 'Reconstitution of Prospermatogonial Specification in Vitro from Human Induced Pluripotent Stem Cells'. Nature Communications 11, no. 1 (9 November 2020): 5656. https://doi.org/10.1038/s41467-020-19350-3.

[63]Tang, Walfred W. C., Toshihiro Kobayashi, Naoko Irie, Sabine Dietmann, and M. Azim Surani. 'Specification and Epigenetic Programming of the Human Germ Line'. Nature Reviews Genetics 17, no. 10 (October 2016): 585–600. https://doi.org/10.1038/nrg.2016.88.

Various research groups are working on elements of IVO. Commercial ventures aiming to treat infertility include Ovelle, Conception, Ivy Natal, Vitra, Gameto, and Dioseve. See the citations in this section for academic groups.

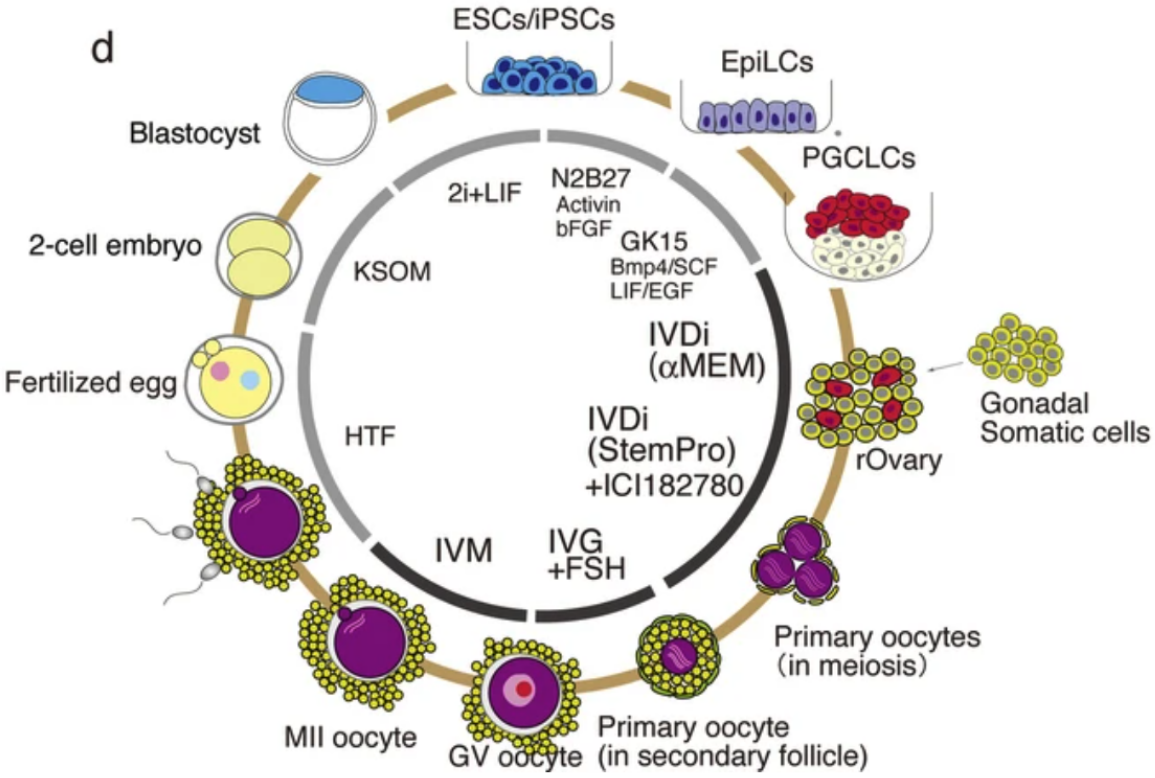
Sneak peek (but note that some of these steps still can't be said to be "solved"):


The path to in vitro oogenesis
somatic cell
Takahashi and Yamanaka, 2006
induced pluripotent stem cell (iPSC)
Irie et al. 2015, Sasaki et al. 2015
primordial germ cell-like cell (PGCLC)
Yamashiro et al. 2020
oogonium-like cell
Mizuta et al. 2022
small dictyate primary oocyte, primordial follicle
???
in vitro differentiation
germinal vesicle oocyte, antral follicle
Silber et al. 2024
mature secondary oocyte, metaphase II arrested
Berkeley Genomics Project
2025 license: CC BY-SA 4.0

{width:

58% }

Morohaku et al. (2016) were able to culture fetal mouse ovaries in vitro ("ex vivo") so that the PGCs in the ovaries grew into oocytes[64]; this method uses whole tissues taken from mouse fetuses. That same year, Hikabe et al. (2016) achieved a kind of end-to-end IVO in mice. They took mouse PGC-like cells created in vitro from ESCs, and grew them in culture with cells taken from female mouse embryo ovaries and with certain culture media. Some of the PGC-like cells developed into oocytes, and when fertilized, those oocytes produced apparently healthy offspring who were themselves naturally fertile[65]. An illustration of their method:



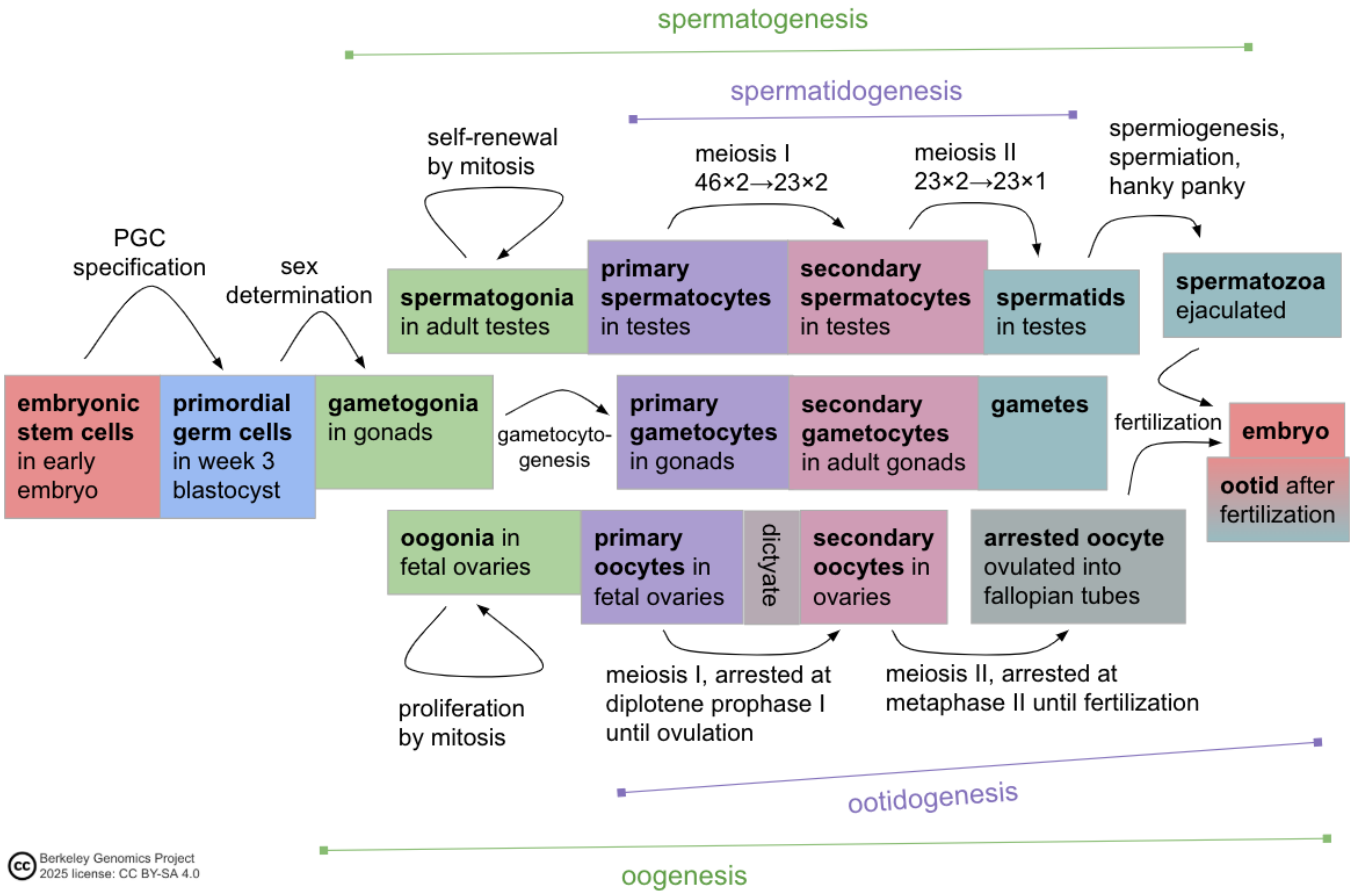
(Extended Data Figure 9d from Hikabe et al. (2016).)

The human analog of this method would not be very practical for use in reproductive GV, and mouse reproductive epigenetics is different from human reproductive epigenetics, but it does demonstrate the possibility of mammalian IVO.

When it comes to humans, there's been significant progress on several legs of the female germline's developmental journey, though full IVO has not been achieved. Recall the process of natural oogenesis, depicted running along the bottom half of this diagram:

[64]Morohaku, Kanako, Ren Tanimoto, Keisuke Sasaki, Ryouka Kawahara-Miki, Tomohiro Kono, Katsuhiko Hayashi, Yuji Hirao, and Yayoi Obata. 'Complete in Vitro Generation of Fertile Oocytes from Mouse Primordial Germ Cells'. Proceedings of the National Academy of Sciences of the United States of America 113, no. 32 (9 August 2016): 9021–26. https://doi.org/10.1073/pnas.1603817113.

[65]Hikabe, Orie, Nobuhiko Hamazaki, Go Nagamatsu, Yayoi Obata, Yuji Hirao, Norio Hamada, So Shimamoto, et al. 'Reconstitution in Vitro of the Entire Cycle of the Mouse Female Germ Line'. Nature 539, no. 7628 (November 2016): 299–303. https://doi.org/10.1038/nature20104.



As mentioned above, we can create pluripotent stem cells from human somatic cells (iPSCs)[66], and from iPSCs we can create human PGC-like cells[6768]. The next step is to take PGC-like cells and turn them into oogonia, the stem cells that proliferate in the fetal ovaries and then differentiate into oocytes (or die).

Yamashiro et al. (2020) produced human oogonia-like cells from human iPSCs. They induced the iPSCs into PGC-like cells, and then cultured the iPGCLCs with mouse fetal ovary cells, producing oogonia-like cells[69]. Their method takes around 2.5 months, and around 10% of starting hiPGCLCs make it to an oogonia-like stage. As they suggest, mouse fetal ovarian cells probably won't be a good platform for human oogenesis proper, because they probably can't support the differentiation of oogonia into meiotic oocytes, especially given the close interaction between a natural developing oocyte and its surrounding ovarian support cells.

The next step from oogonia is oogenesis, where the oogonium differentiates into a primary oocyte, which (eventually) goes through meiosis I to become a secondary oocyte, and then meiosis II to become an ovum. An illustration:

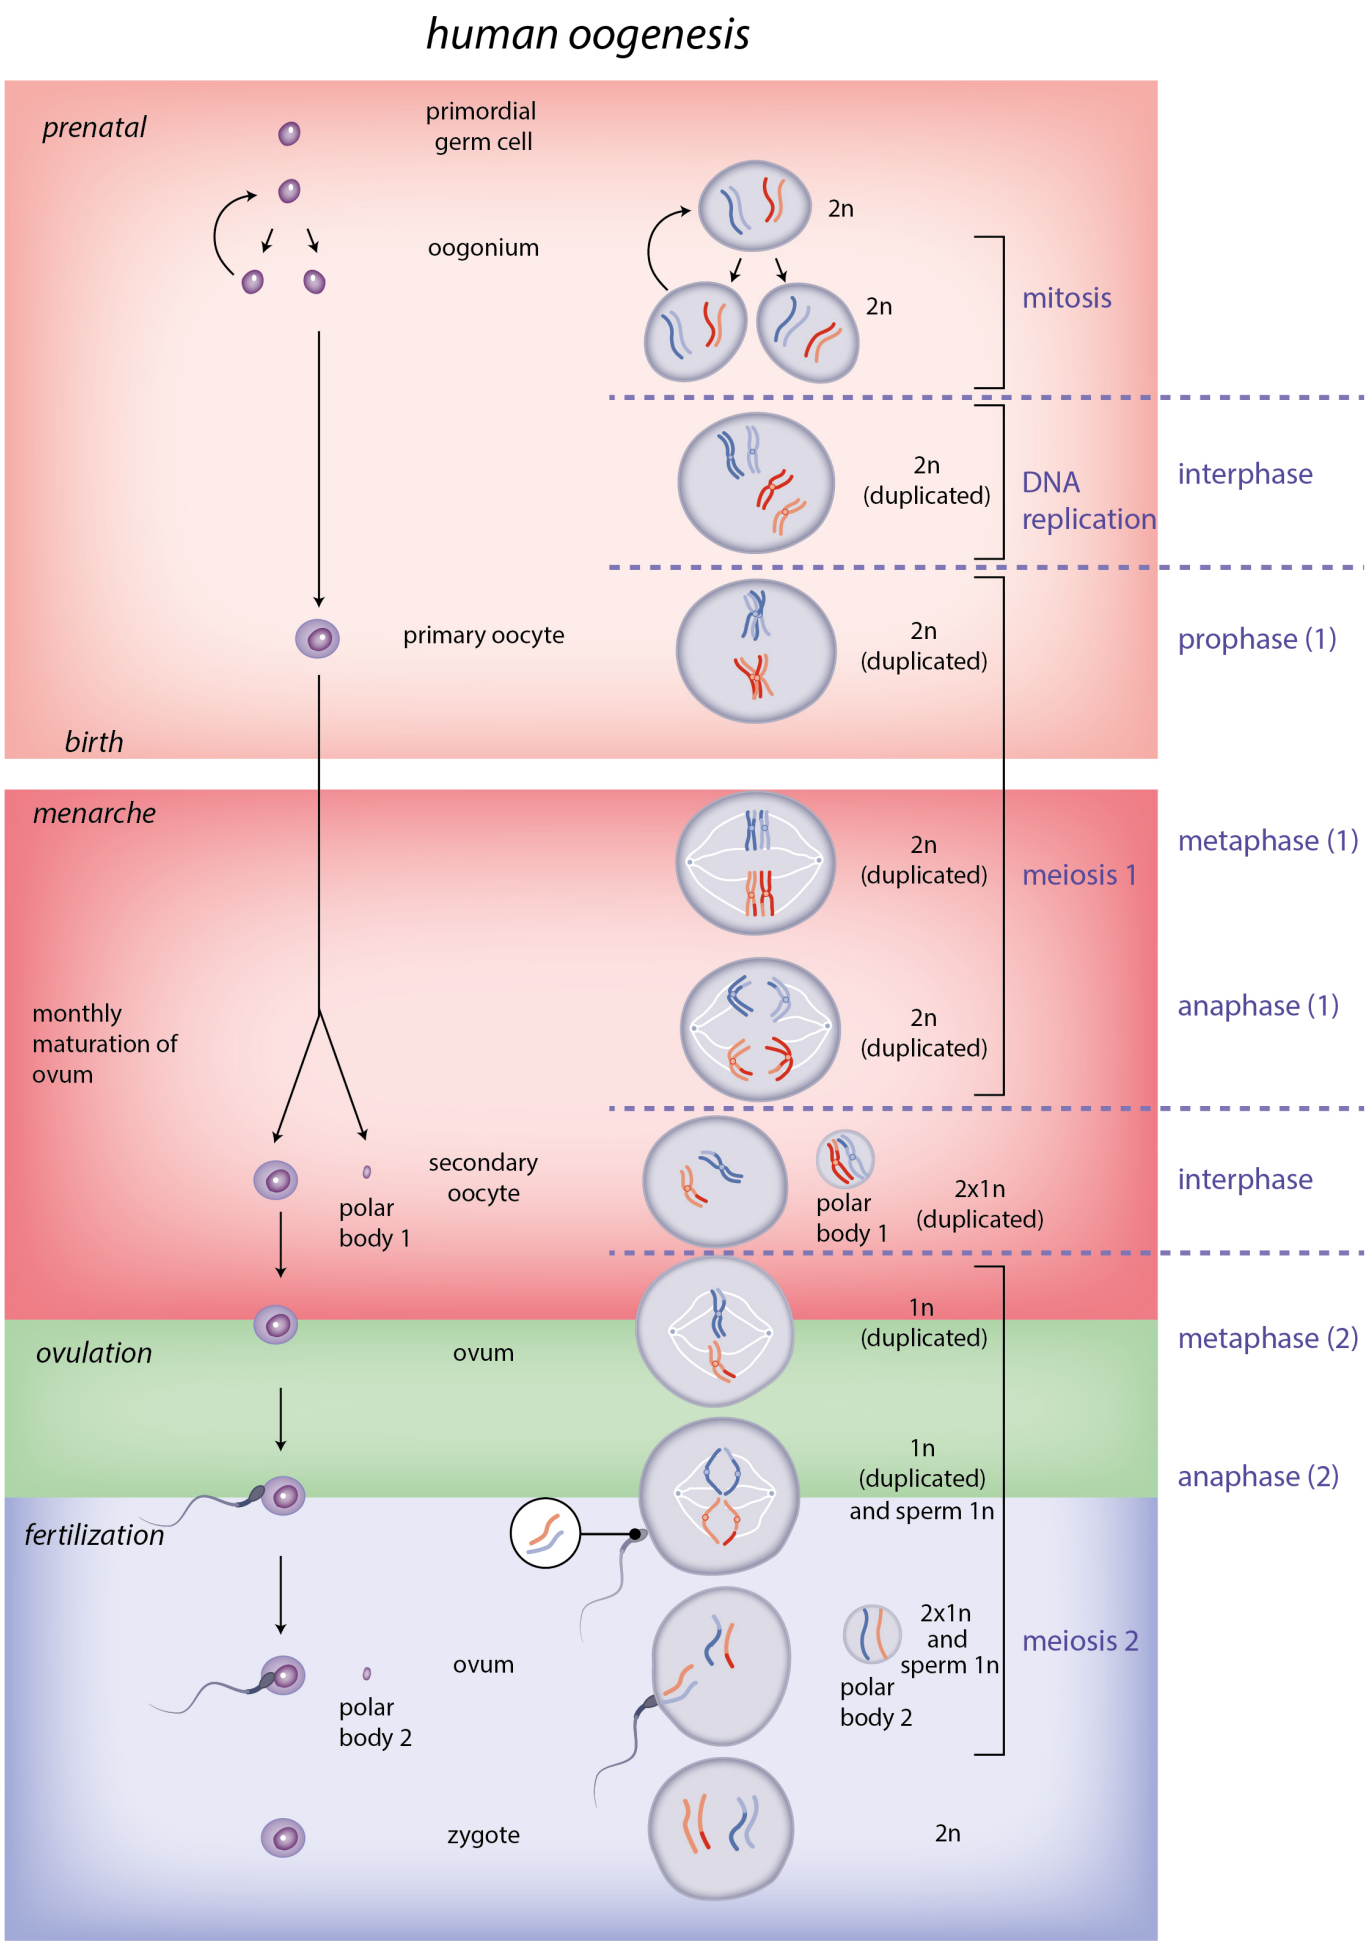
---

[66]Takahashi, Kazutoshi, and Shinya Yamanaka. 'Induction of Pluripotent Stem Cells from Mouse Embryonic and Adult Fibroblast Cultures by Defined Factors'. Cell 126, no. 4 (25 August 2006): 663–76. https://doi.org/10.1016/j.cell.2006.07.024.

[67]Irie, Naoko, Leehee Weinberger, Walfred W.C. Tang, Toshihiro Kobayashi, Sergey Viukov, Yair S. Manor, Sabine Dietmann, Jacob H. Hanna, and M. Azim Surani. 'SOX17 Is a Critical Specifier of Human Primordial Germ Cell Fate'. Cell 160, no. 1–2 (15 January 2015): 253–68. https://doi.org/10.1016/j.cell.2014.12.013.

[68]Sasaki, Kotaro, Shihori Yokobayashi, Tomonori Nakamura, Ikuhiro Okamoto, Yukihiro Yabuta, Kazuki Kurimoto, Hiroshi Ohta, et al. 'Robust In Vitro Induction of Human Germ Cell Fate from Pluripotent Stem Cells'. Cell Stem Cell 17, no. 2 (6 August 2015): 178–94. https://doi.org/10.1016/j.stem.2015.06.014.

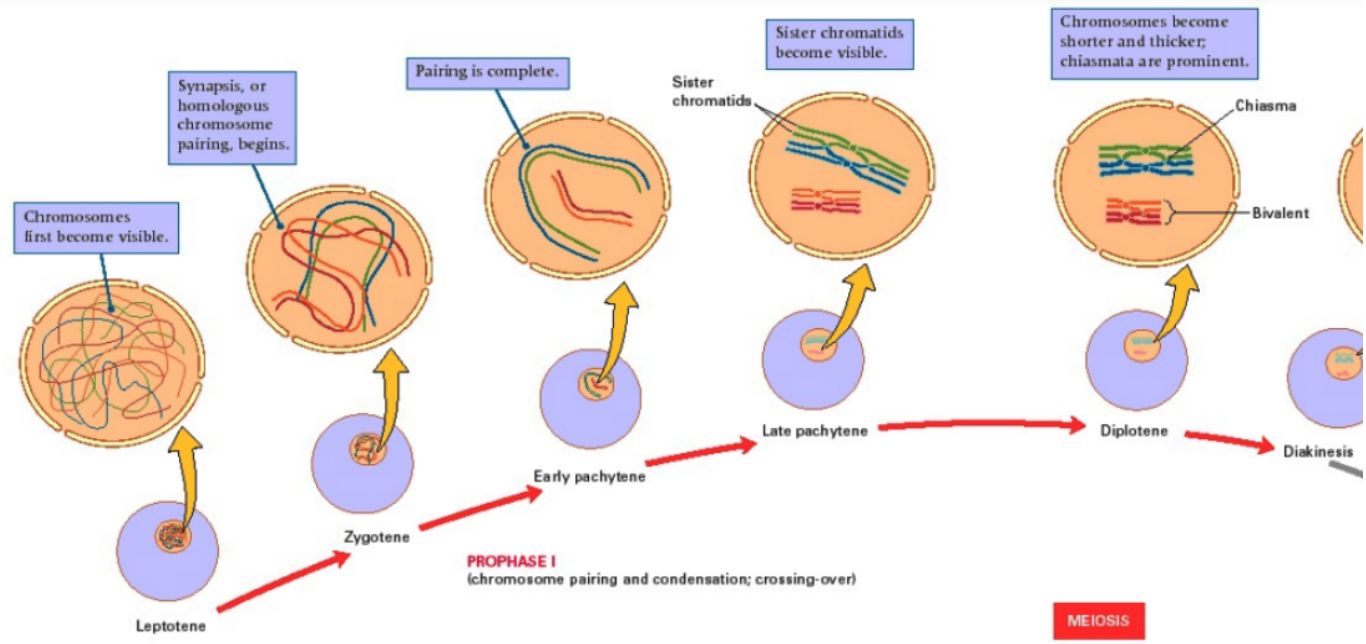
[69]Yamashiro, Chika, Kotaro Sasaki, Shihori Yokobayashi, Yoji Kojima, and Mitinori Saitou. 'Generation of Human Oogonia from Induced Pluripotent Stem Cells in Culture'. Nature Protocols 15, no. 4 (April 2020): 1560–83. https://doi.org/10.1038/s41596-020-0297-5.



(Diagram from Slagter et al.[70].)

[70]Ron Slagter, O. Paul Gobée, LUMC, Hope Wicks, LUMC, et al. 'Slagter - Drawing Human Oogenesis Diagram - English Labels | AnatomyTOOL'. Accessed 21 February 2025. https://anatomytool.org/content/slagter-drawing-human-oogenesis-diagram-english-labels.

The first step is to enter meiosis I. In contrast to spermatogenesis, in oogenesis the primary gametocyte, in this case a primary oocyte, does not complete meiosis I at first. See the gap shown in the above diagram between birth and menarche; in that gap, meiosis is paused. Primary oocytes get most of the way through prophase I of meiosis I, and then stop in the diplotene stage. Most of the meiotic action has already happened at that point. The genome has been duplicated. Homologous chromosomes have paired up and have formed chiasmata, which become DNA recombination sites, where DNA is exchanged (or in small amounts copied) between homologous chromosomes. Zooming in to prophase I on the earlier diagram of meiosis:



As mentioned above, Smela et al. (2024) induced human iPSCs to get to the pachytene substage of prophase I, and some cells looked like fully meiotic cells in terms of their RNA expression[71]. Their method somewhat bypasses the PGC and oogonium stages (though there was some resemblance); that might make it harder to produce epigenomically correct oocytes, but on the other hand it's cheaper and faster. Further, their method doesn't use extracted tissue, making it much more scalable. Instead, they use culture media and gene overexpression. They knew which genes to try overexpressing by looking at Garcia-Alonso et al. (2022)'s human gonadal tissue dataset of single-cell RNA-seq (which says what RNAs are in a cell, hence what genes are being expressed) and scATAC-seq (which says which areas of a chromosome are packed/unpacked, hence likely to be inactive/active)[72].

Mizuta et al. (2022) developed methods for taking ovarian tissue from cynomolgus macaque fetuses and culturing them. These reconstituted ovaries could make cyno cells to advance to the diplotene substage of prophase I, according to transcriptomics (sequencing mRNAs). (I'm not sure whether they used cyno iPSCs or just monitored the oogonia in their reconstituted ovaries.) They got similar results with human reconstituted ovaries, where the oogonia present in the extracted fetal ovarian tissue developed to what appeared to be the diplotene stage. Since the method uses tissue from aborted fetuses, it can't scale, but it's a step[73]. No one has made oogonia go all the way through meiosis to produce a mature oocyte.

Natural oogenesis is *not* a process where isolated cells, destined to be eggs, go through development on their own. Rather, oogonia are connected to each other during proliferation. As the oogonia differentiate into oocytes, they recruit other cells from the ovaries to form follicles, which are support structures that grow with and coregulate the oocyte. Follicles are tightly connected to oocytes: the cumulus granulosa cells, which directly border the oocyte, literally form direct cytoplasmic bridges with the oocyte[74]. Follicles are

[71] Pierson Smela, Merrick, Jessica Adams, Carl Ma, Laura Breimann, Ursula Widocki, Toshi Shioda, and George M. Church. 'Induction of Meiosis from Human Pluripotent Stem Cells'. bioRxiv, 31 May 2024, 2024.05.31.596483. https://doi.org/10.1101/2024.05.31.596483.

[72] Garcia-Alonso, Luz, Valentina Lorenzi, Cecilia Icoresi Mazzeo, João Pedro Alves-Lopes, Kenny Roberts, Carmen Sancho-Serra, Justin Engelbert, et al. 'Single-Cell Roadmap of Human Gonadal Development'. Nature 607, no. 7919 (July 2022): 540–47. https://doi.org/10.1038/s41586-022-04918-4.

[73] Mizuta, Ken, Yoshitaka Katou, Baku Nakakita, Aoi Kishine, Yoshiaki Nosaka, Saki Saito, Chizuru Iwatani, et al. 'Ex Vivo Reconstitution of Fetal Oocyte Development in Humans and Cynomolgus Monkeys'. The EMBO Journal 41, no. 18 (15 September 2022): e110815. https://doi.org/10.15252/embj.2022110815.

[74] Baena, Valentina, and Mark Terasaki. 'Three-Dimensional Organization of Transzonal Projections and Other Cytoplasmic Extensions in the Mouse Ovarian Follicle'. Scientific Reports 9, no. 1 (4 February 2019): 1262. https://doi.org/10.1038/s41598-

crucial for supplying nutrients to the growing oocytes, and for regulating the meiosis and reprogramming processes that the oocyte undergoes. Here's a diagram depicting human oogenesis and folliculogenesis:

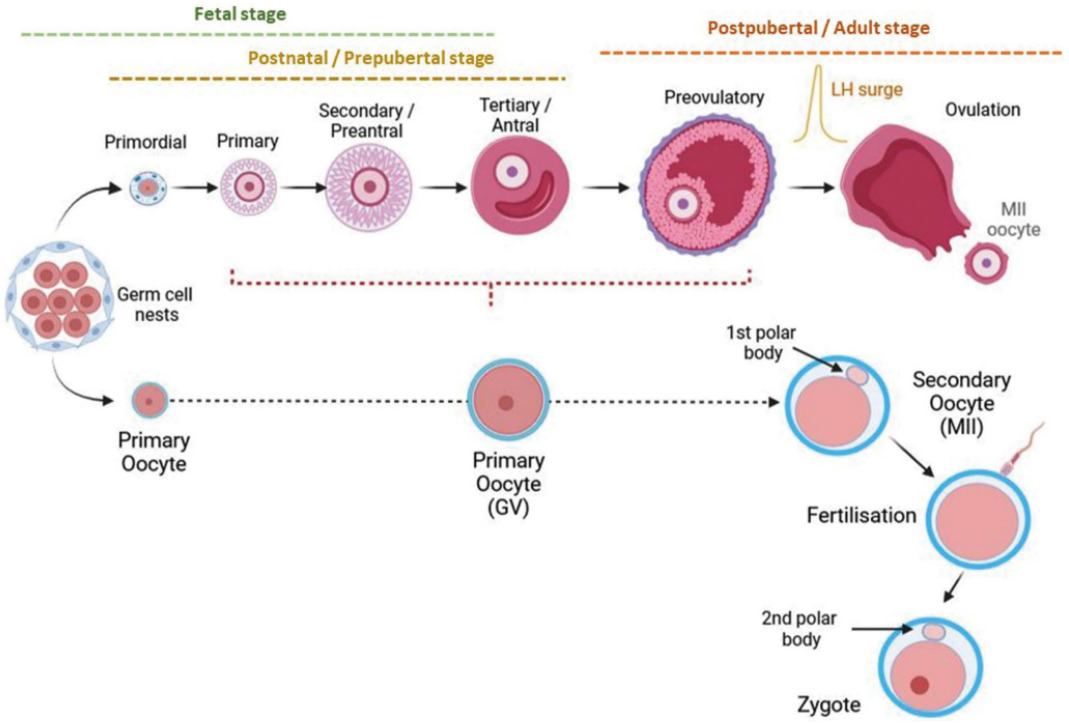


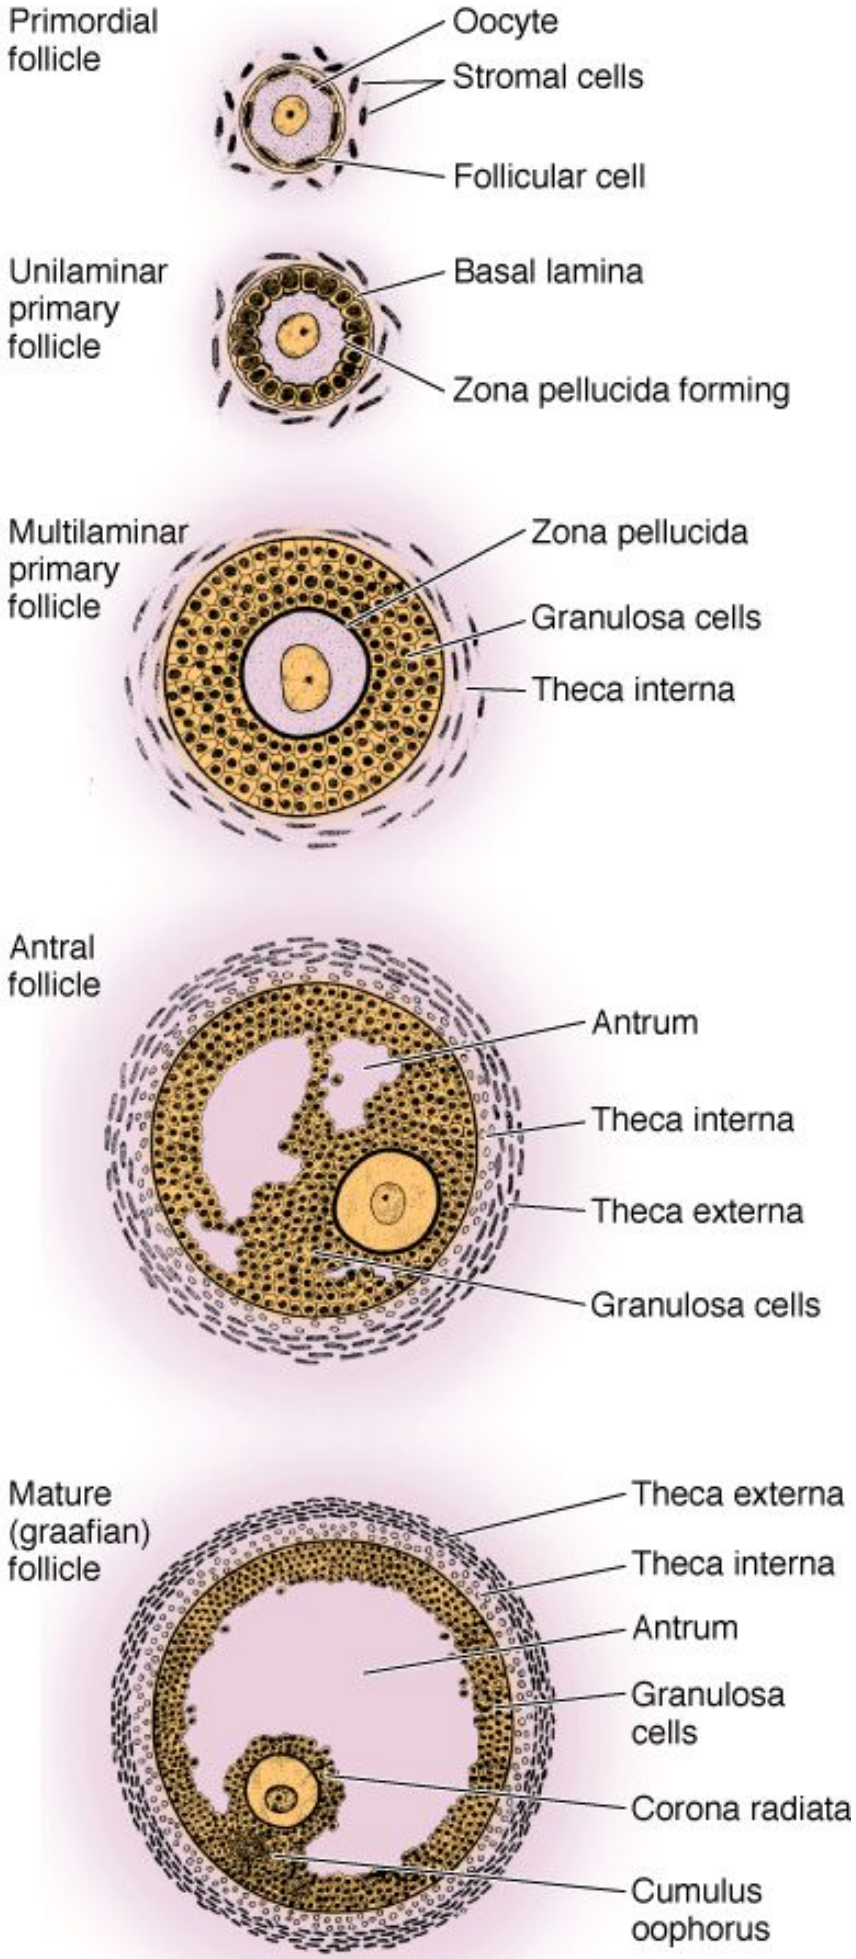
FIGURE 11.1 Schematic representation of the simultaneous events that encompass folliculogenesis and oogenesis. Early follicle formation takes place during fetal development with folliculogenesis progressing throughout female life, but final follicle growth and ovulation only occur after puberty, with the preovulatory LH surge. Diagram created using BioRender.com.

(Figure 11.1 from chapter 11 of Campell and Maalouf 2024[75].)

More zoomed in:

018-37766-2.

[75]Campbell, Alison, and Walid Maalouf. Mastering Clinical Embryology: Good Practice, Clinical Biology, Assisted Reproductive Technologies, and Advanced Laboratory Skills. CRC Press, 2024, https://www.google.com/books/edition/Mastering_Clinical_Embryology/xkb3EAAAQBAJ.



Source: Mescher AL: *Junqueira's Basic Histology: Text and Atlas, 12th Edition:* http://www.accessmedicine.com

Copyright © The McGraw-Hill Companies, Inc. All rights reserved.

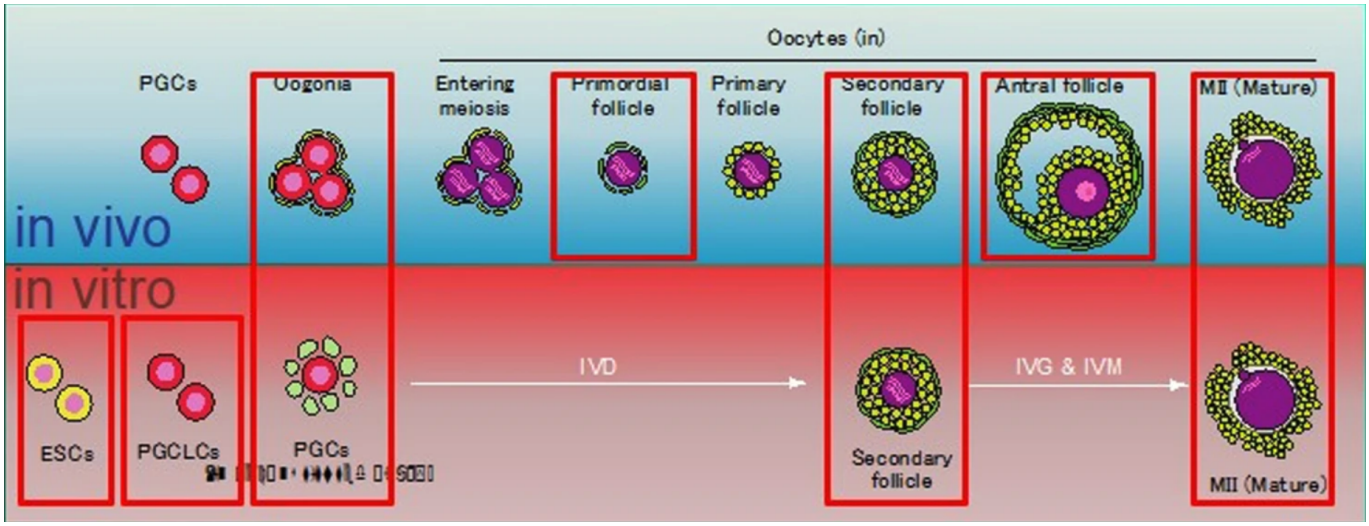
{width: 45% }

(Diagram from Mescher (2025)[76].)

The natural process of postnatal, preovulatory oogenesis and folliculogenesis is somewhat murky. Primordial follicles are formed during fetal development and stay stored in the ovaries, with the oocytes arrested in the diplotene substage of meiosis I. But more developed follicles, through the preantral and antral stages of growth, are found in postnatal prepubertal ovaries[77]. On the other hand, prepubertal ovaries are not the same as pubertal ovaries, even leaving aside ovulation; follicles in prepubertal ovaries grow at much lower rates in culture compared to pubertal ovaries, suggesting that follicles and oocytes may continue with important development during all of childhood[78]. Unless this was clarified recently, it's not well-understood what development happens in follicles and oocytes during childhood, so we don't know what it would take to recapitulate that development in vitro, and we don't know how necessary it is to do so in order to make reproductively competent oocytes in vitro.

Thus, there's a gap. We can sort of get from a human somatic cell to a human primary oocyte-like cell, albeit inefficiently and expensively, but as yet we haven't proceeded from that point.

What comes after the gap? The next step would be the final steps of oogenesis: maturation of the follicle and oocyte past the antral stage, partway through meiosis II, at which point the secondary oocyte is ready to be fertilized. This maturation normally happens during ovulation. So to recap, there's the gap, where in vitro differentiation should go, taking primordial follicles to secondary or antral follicles and germinal vesicle oocytes; and then there's in vitro maturation, which takes the grown follicles and produces mature oocytes from them. An illustration:



(Figure 1 from Silber et al. (2024)[79].)

It's possible to preserve a woman's fertility, e.g. through cancer treatments, by extracting a part of her ovaries. The ovarian tissue contains oocytes in follicles that are in a variety of growth stages. It's then possible to get the oocytes and follicles to mature. They grow and meiose into secondary oocytes ready to be fertilized, producing live human births[80][81][82]. Silber et al. (2024) describe how, basically, squeezing the ovarian tissue produces matured oocytes with high efficiency[83].

[76] Mescher, Anthony L. 'Junqueira's Basic Histology: Text and Atlas, 17th Edition | AccessMedicine | McGraw Hill Medical'. Accessed 21 February 2025. https://accessmedicine.mhmedical.com/book.aspx?bookID=3390.

[77] Peters, H., A. G. Byskov, and J. Grinsted. 'Follicular Growth in Fetal and Prepubertal Ovaries of Humans and Other Primates'. Clinics in Endocrinology and Metabolism 7, no. 3 (November 1978): 469–85. https://doi.org/10.1016/s0300-595x(78)80005-x.

[78] Anderson, R.A., M. McLaughlin, W.H.B. Wallace, D.F. Albertini, and E.E. Telfer. 'The Immature Human Ovary Shows Loss of Abnormal Follicles and Increasing Follicle Developmental Competence through Childhood and Adolescence'. Human Reproduction (Oxford, England) 29, no. 1 (January 2014): 97–106. https://doi.org/10.1093/humrep/det388.

[79] Silber, Sherman J., Sierra Goldsmith, Leilani Castleman, and Katsuhiko Hayashi. 'In Vitro Maturation, In Vitro Oogenesis, and Ovarian Longevity'. Reproductive Sciences 31, no. 5 (1 May 2024): 1234–45. https://doi.org/10.1007/s43032-023-01427-1.

[80] De Vos, Michel, Michaël Grynberg, Tuong M. Ho, Ye Yuan, David F. Albertini, and Robert B. Gilchrist. 'Perspectives on the Development and Future of Oocyte IVM in Clinical Practice'. Journal of Assisted Reproduction and Genetics 38, no. 6 (June 2021): 1265–80. https://doi.org/10.1007/s10815-021-02263-5.

[81] Kawamura, Kazuhiro, Yuan Cheng, Nao Suzuki, Masashi Deguchi, Yorino Sato, Seido Takae, Chi-hong Ho, et al. 'Hippo Signaling Disruption and Akt Stimulation of Ovarian Follicles for Infertility Treatment'. Proceedings of the National Academy of Sciences 110, no. 43 (22 October 2013): 17474–79. https://doi.org/10.1073/pnas.1312830110.

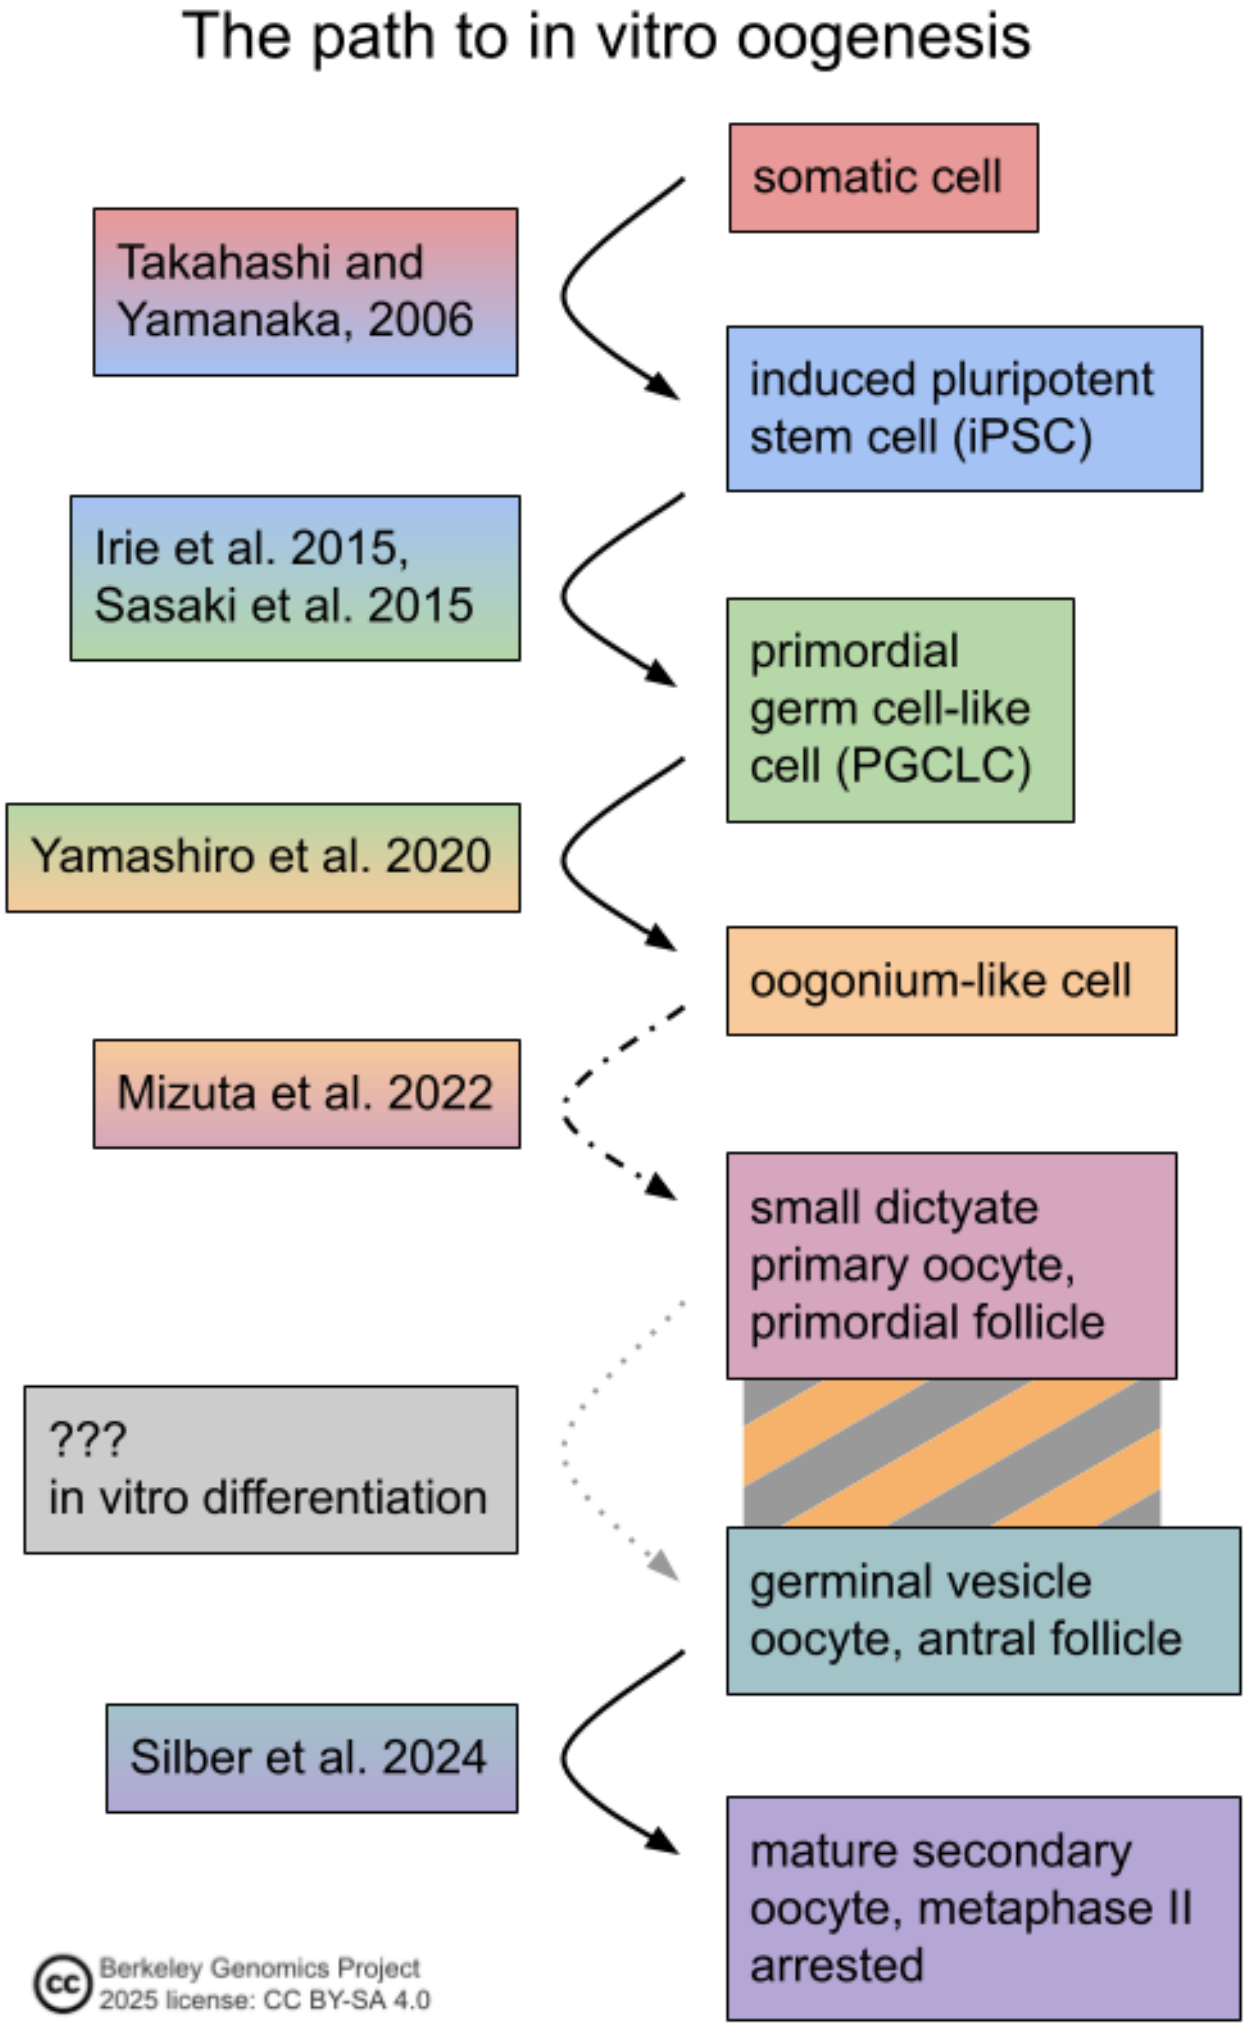
[82] Hatırnaz, Şafak, Barış Ata, Ebru Saynur Hatırnaz, Michael Haim Dahan, Samer Tannus, Justin Tan, and Seang Lin Tan. 'Oocyte in Vitro Maturation: A Sytematic Review'. Turkish Journal of Obstetrics and Gynecology 15, no. 2 (June 2018): 112–25. https://doi.org/10.4274/tjod.23911.

[83] Silber, Sherman J., Sierra Goldsmith, Leilani Castleman, and Katsuhiko Hayashi. 'In Vitro Maturation, In Vitro Oogenesis,

By stitching all these methods together—creating hiPSCs, then hiPGCLCs, then oogonia-like cells, then dictyate-arrested primary oocyte-like cells with primordial follicles, and then ???, and then in vitro maturation—it looks almost possible to make a (very slow, inefficient, complicated) minimum viable human IVO method:

---

and Ovarian Longevity'. Reproductive Sciences 31, no. 5 (1 May 2024): 1234–45. https://doi.org/10.1007/s43032-023-01427-1.


The path to in vitro oogenesis
somatic cell
Takahashi and Yamanaka, 2006
induced pluripotent stem cell (iPSC)
Irie et al. 2015, Sasaki et al. 2015
primordial germ cell-like cell (PGCLC)
Yamashiro et al. 2020
oogonium-like cell
Mizuta et al. 2022
small dictyate primary oocyte, primordial follicle
???
in vitro differentiation
germinal vesicle oocyte, antral follicle
Silber et al. 2024
mature secondary oocyte, metaphase II arrested
Berkeley Genomics Project
2025 license: CC BY-SA 4.0

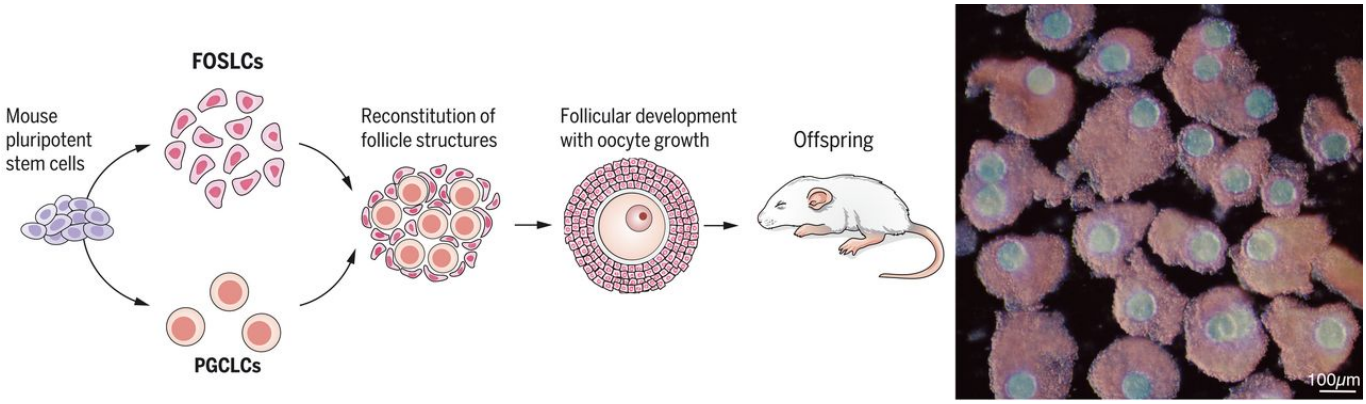
{width:

55% }

But in fact, this gap is probably quite significant: It is believed to be when maternal imprinting is established. This is the crucial oocyte growth period, starting with the primary oocyte recruiting a primordial follicle during fetal development, and extending through some amount of postnatal growth, to childhood or adulthood, producing a grown primary oocyte in a secondary or antral follicle[84][85][86]. Natural oocytes also aren't meiotically competent until toward the end of their pre-antral growth phase, though this may be due to the follicle support cells actively suppressing meiosis[87].

To cross this gap, we'd likely have to make some form of artificial ovarian tissue. Possibly this could be ovarian tissue reconstituted from ovarian cells taken from another mammal, but it seems strange for this to work, given how closely granulosa cells interface with the oocyte. Another possibility would be to use human adult ovarian tissue, e.g. from the mother looking to have a child, but this is quite costly, especially since it involves surgery. I also don't expect that it would work; adult ovaries are not the same as postnatal ovaries. More research is needed into growing follicles in vitro[88].

Another approach is to induce stem cells to differentiate into ovarian-like cells, maybe forming 3D organoids. Smela et al. (2023) differentiated human iPSCs into cells that resemble granulosa cells, a key component of follicles[89]. Yoshino et al. (2021) made a functional oocyte by inducing pluripotent stem cells to differentiate into fetal ovarian somatic cell-like cells which could support oogenesis… *(say it with me)*… in mice[90]:



(Diagram from Yoshino et al. (2021).)

In vitro spermatogenesis (IVS)

Given more research, it would also be possible to create sperm from stem cells in vitro—i.e., do in vitro spermatogenesis (IVS). This possibility constitutes a path to epigenomic correction that's somewhat independent of IVO; we could plausibly be able to do IVS and not IVO, or vice versa.

Takeaways:

- IVS research, taken together, has already found ways to derive haploid cells from diploid cells through meiosis in vitro.

[84]Saitou, Mitinori, and Hidetaka Miyauchi. 'Gametogenesis from Pluripotent Stem Cells'. Cell Stem Cell 18, no. 6 (June 2016): 721–35. https://doi.org/10.1016/j.stem.2016.05.001.

[85]Trasler, Jacquetta M. 'Gamete Imprinting: Setting Epigenetic Patterns for the next Generation'. Reproduction, Fertility and Development 18, no. 2 (2006): 63. https://www.publish.csiro.au/rd/Fulltext/RD05118.

[86]Kelsey, Gavin, and Robert Feil. 'New Insights into Establishment and Maintenance of DNA Methylation Imprints in Mammals'. Philosophical Transactions of the Royal Society B: Biological Sciences 368, no. 1609 (5 January 2013): 20110336. https://doi.org/10.1098/rstb.2011.0336.

[87]Guglielmo, Maria Cristina, and David F. Albertini. 'The Structural Basis for Coordinating Oogenesis and Folliculogenesis'. In Oogenesis, edited by Giovanni Coticchio, David F. Albertini, and Lucia De Santis, 63–73. London: Springer, 2013. https://doi.org/10.1007/978-0-85729-826-3_5.

[88]Malo, Clara, Sara Oliván, Ignacio Ochoa, and Ariella Shikanov. 'In Vitro Growth of Human Follicles: Current and Future Perspectives'. International Journal of Molecular Sciences 25, no. 3 (26 January 2024): 1510. https://doi.org/10.3390/ijms25031510.

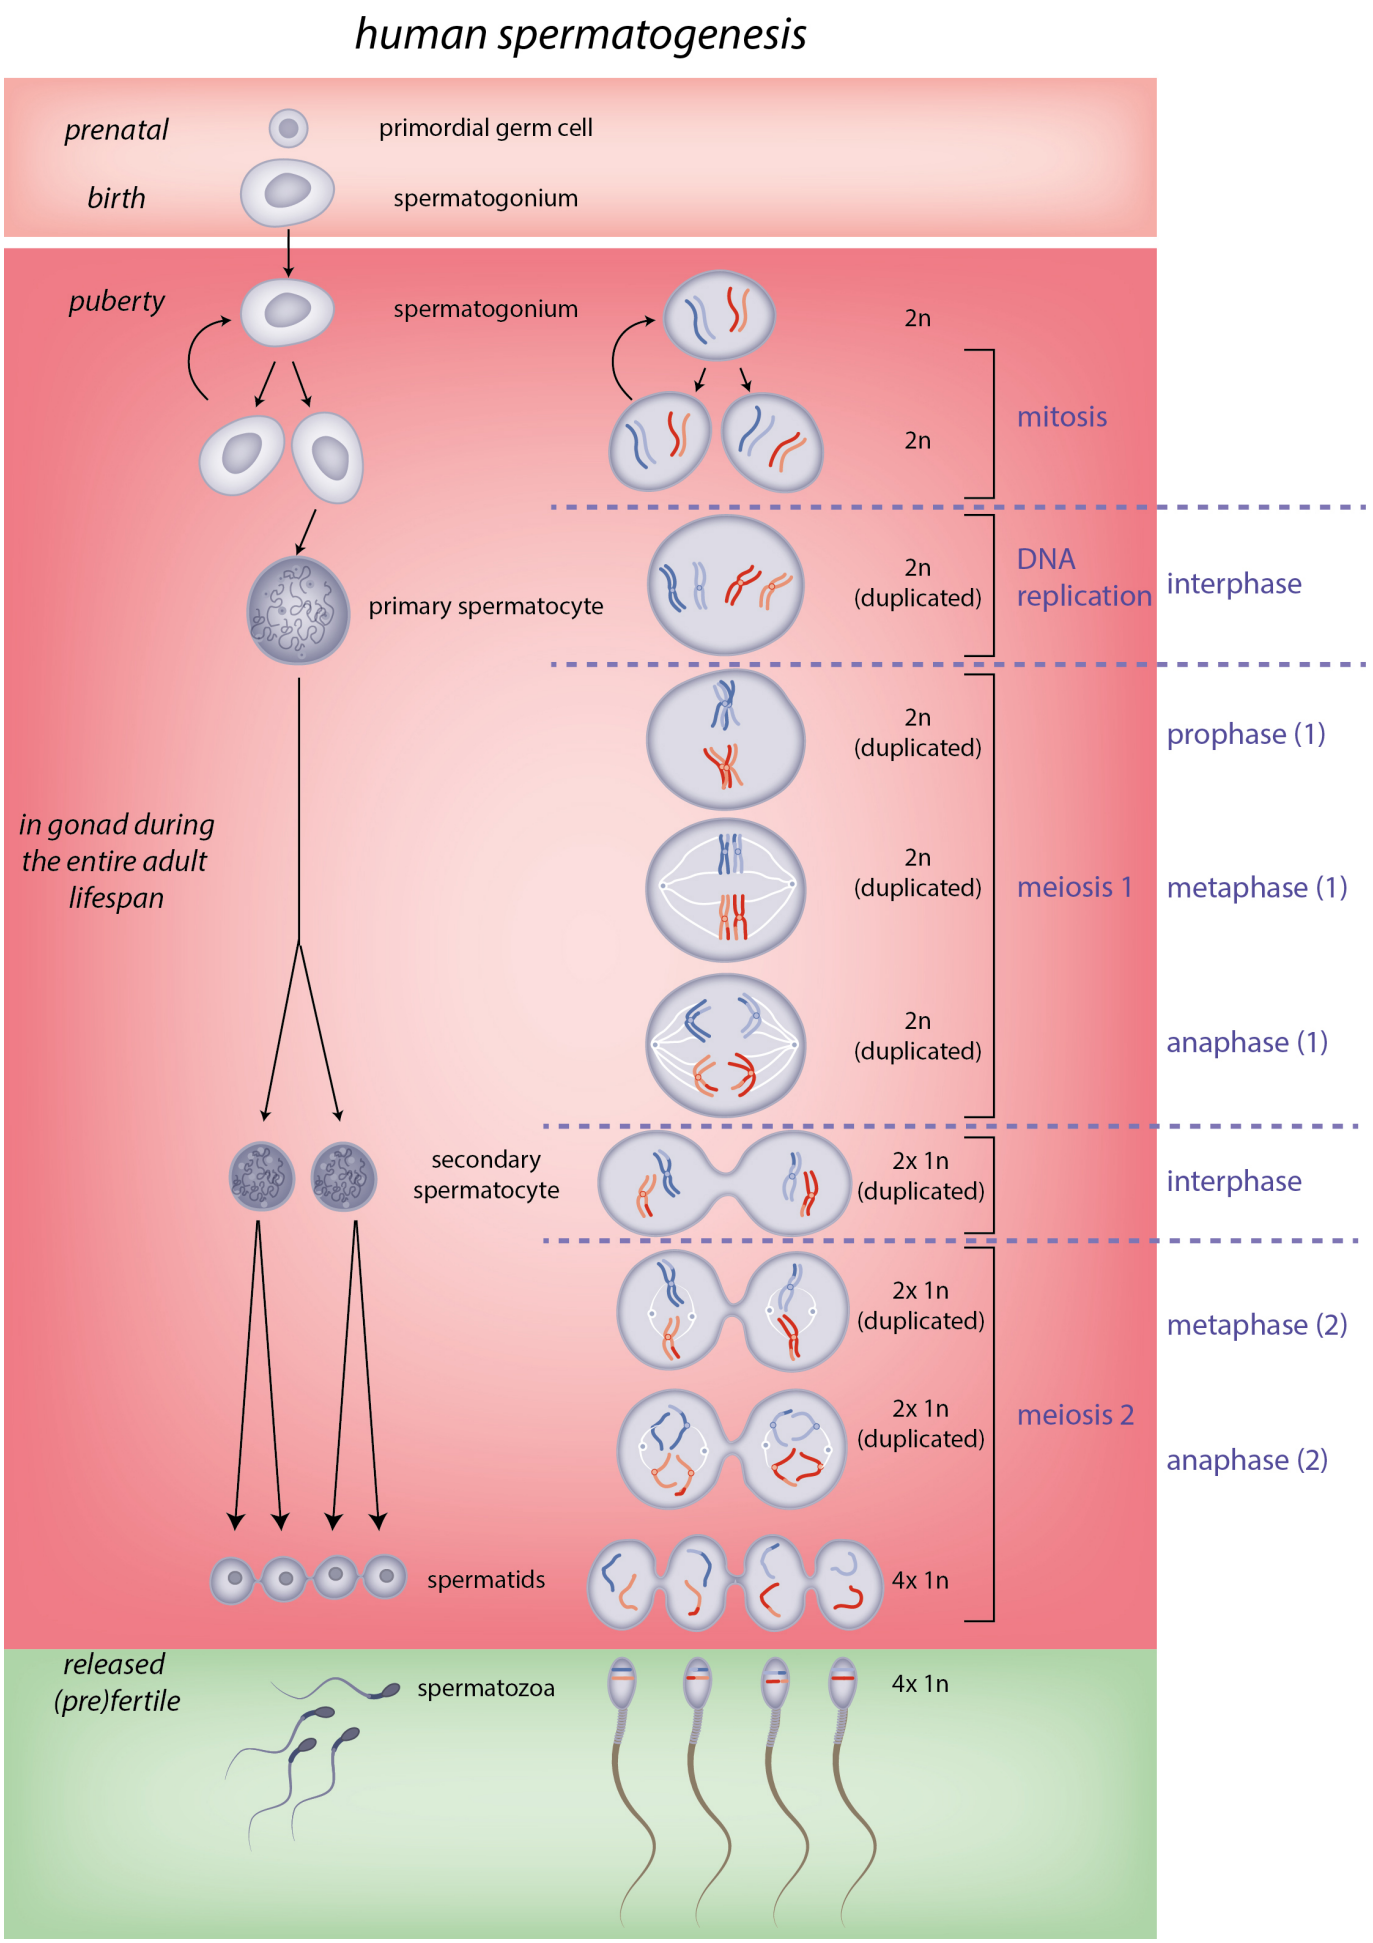
[89]Pierson Smela, Merrick D, Christian C Kramme, Patrick RJ Fortuna, Jessica L Adams, Rui Su, Edward Dong, Mutsumi Kobayashi, et al. 'Directed Differentiation of Human iPSCs to Functional Ovarian Granulosa-like Cells via Transcription Factor Overexpression'. Edited by T Rajendra Kumar and Marianne E Bronner. eLife 12 (21 February 2023): e83291. https://doi.org/10.7554/eLife.83291.

[90]Yoshino, Takashi, Takahiro Suzuki, Go Nagamatsu, Haruka Yabukami, Mika Ikegaya, Mami Kishima, Haruka Kita, et al. 'Generation of Ovarian Follicles from Mouse Pluripotent Stem Cells'. Science (New York, N.Y.) 373, no. 6552 (16 July 2021): eabe0237. https://doi.org/10.1126/science.abe0237.

- Many methods that have been explored are not scalable because they use tissue taken from human fetuses or children. Some methods use adult tissue, which is potentially feasible for adults who are very motivated to use reproductive genomic vectoring, but is probably too costly to support iterative research or GV methods that have to produce many haploids. Methods that use mouse or monkey cells seem promising to explore more.
- Until recently, the results of these studies were not well characterized epigenomically. We haven't had a clear picture of what various stages of the male germline should look like epigenomically, and the cells produced in these experiments haven't been measured epigenomically. This is because single-cell RNA sequencing and large-scale bisulfite methylation sequencing have only recently become widespread, and because single-cell bisulfite sequencing is low-quality. Thus, we don't know exactly how well we're doing, but we can now find out.

Various research groups are working on elements of IVS. One commercial venture aiming to treat male infertility is Paterna. There was another about a decade ago, Kallistem, but they don't seem to be active. See the citations in this section for academic groups.

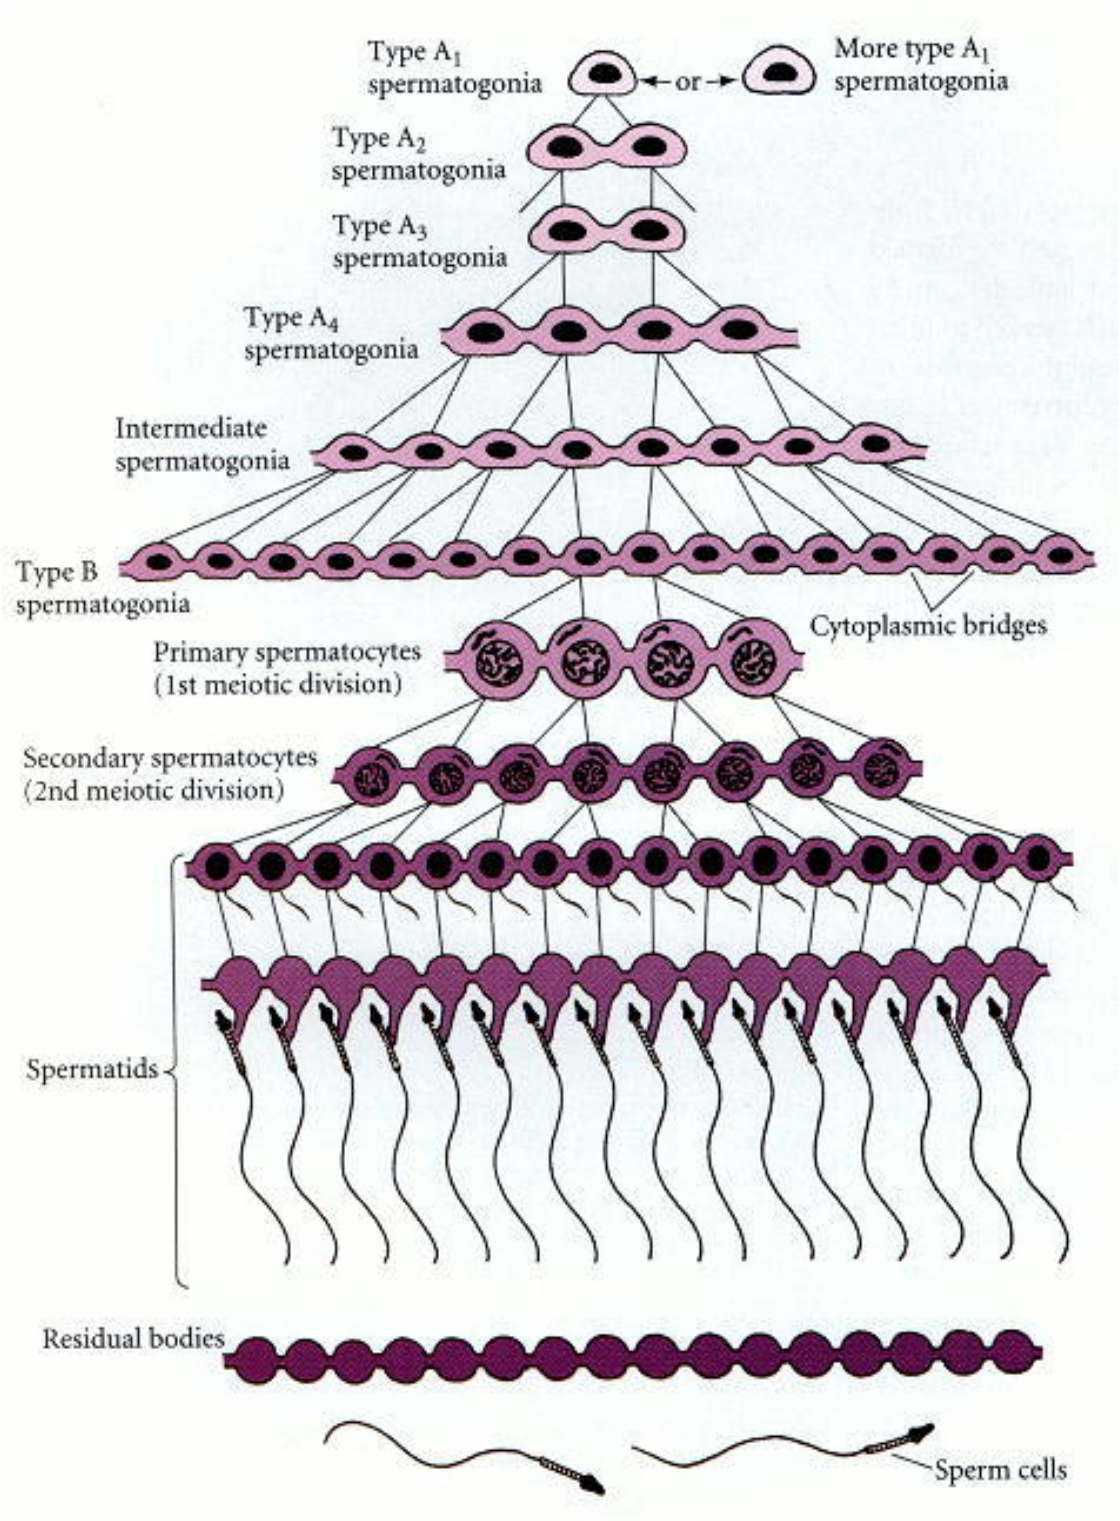
An illustration of spermatogenesis:



(Diagram from Slagter et al.[91].)

[91] Ron Slagter, O. Paul Gobée, LUMC, Hope Wicks, LUMC, et al. 'Slagter - Drawing Human Oogenesis Diagram - English Labels | AnatomyTOOL'. Accessed 23 February 2025. https://anatomytool.org/content/slagter-drawing-human-spermatogenesis-diagram-english-labels.

Unlike oogonia, spermatogonia are active in the adult gonads, spinning off differentiating spermatogonia which will become spermatocytes. Like oogonia, though, the proliferation and differentiation that happens in spermatocytogenesis involves large groups of differentiating spermatogonia that remain connected to each other. In fact, they stay connected all the way through meiosis (unlike primary oocytes which have separated before meiosis):



(Figure 19.18 from Gilbert (2000)[92].)

IVS has been (minimally) achieved in mice[93][94] and rats[95]. Ishikura et al. (2021) reconstituted the whole

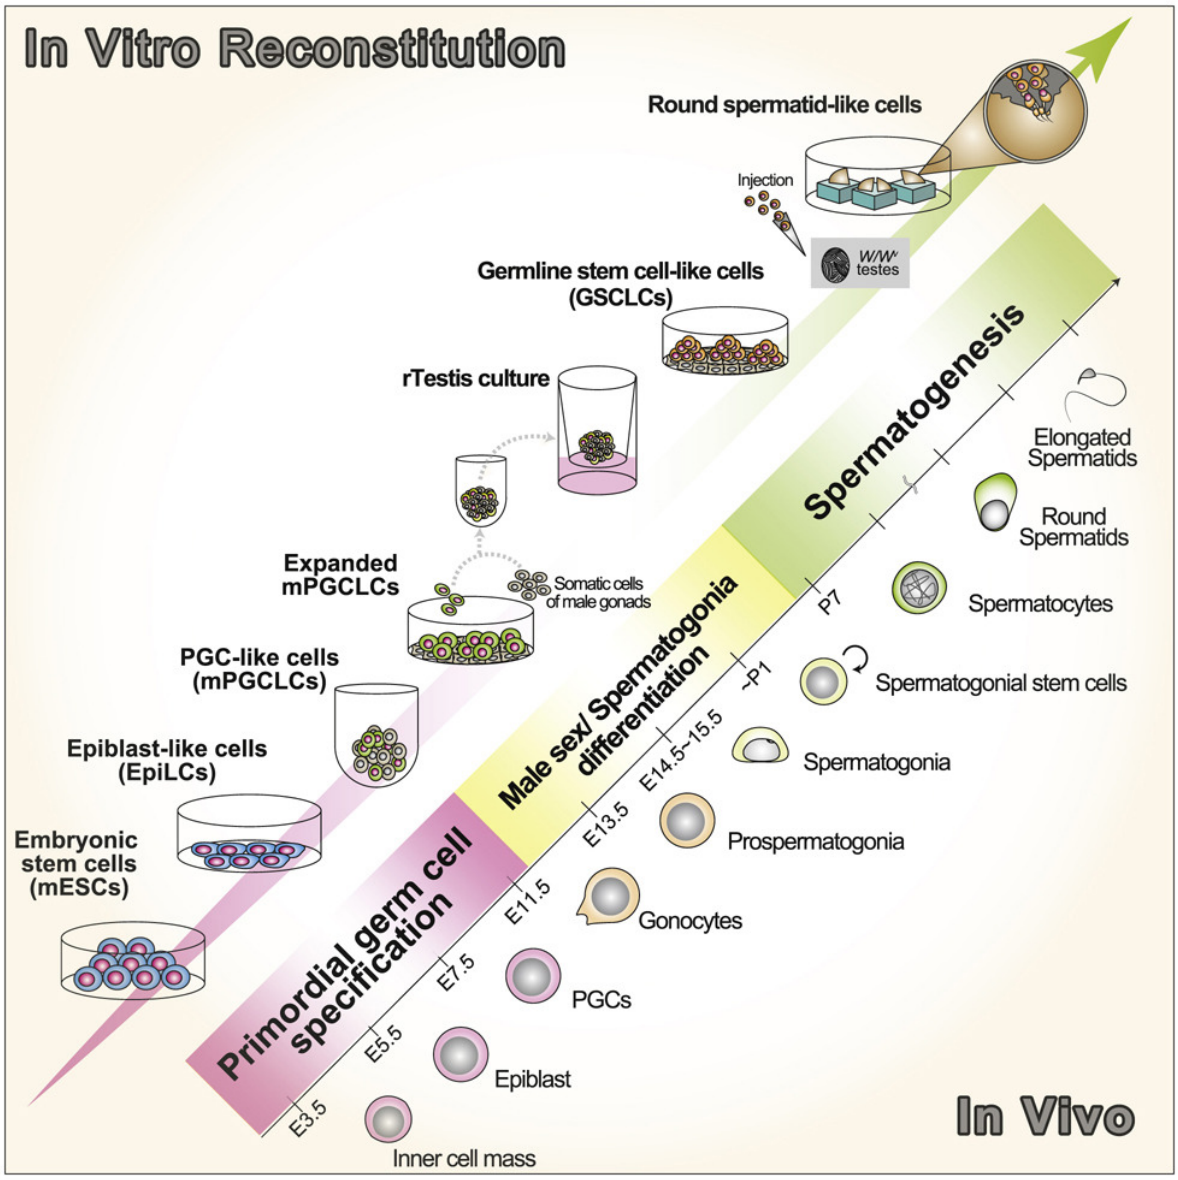
[92]Gilbert, Scott F. 'Spermatogenesis'. In Developmental Biology. 6th Edition. Sinauer Associates, 2000. https://www.ncbi.nlm.nih.gov/books/NBK10095/.

[93]Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.

[94]Sato, Takuya, Kumiko Katagiri, Ayako Gohbara, Kimiko Inoue, Narumi Ogonuki, Atsuo Ogura, Yoshinobu Kubota, and Takehiko Ogawa. 'In Vitro Production of Functional Sperm in Cultured Neonatal Mouse Testes'. Nature 471, no. 7339 (24 March 2011): 504–7. https://doi.org/10.1038/nature09850.

[95]Matsumura, Takafumi, Kumiko Katagiri, Tatsuma Yao, Yu Ishikawa-Yamauchi, Shino Nagata, Kiyoshi Hashimoto, Takuya Sato, et al. 'Generation of Rat Offspring Using Spermatids Produced through in Vitro Spermatogenesis'. Scientific Reports 13,

mouse germline process in vitro[96]. They used reconstituted mouse testes, i.e. blobs of cells from mouse fetal gonadal tissue:

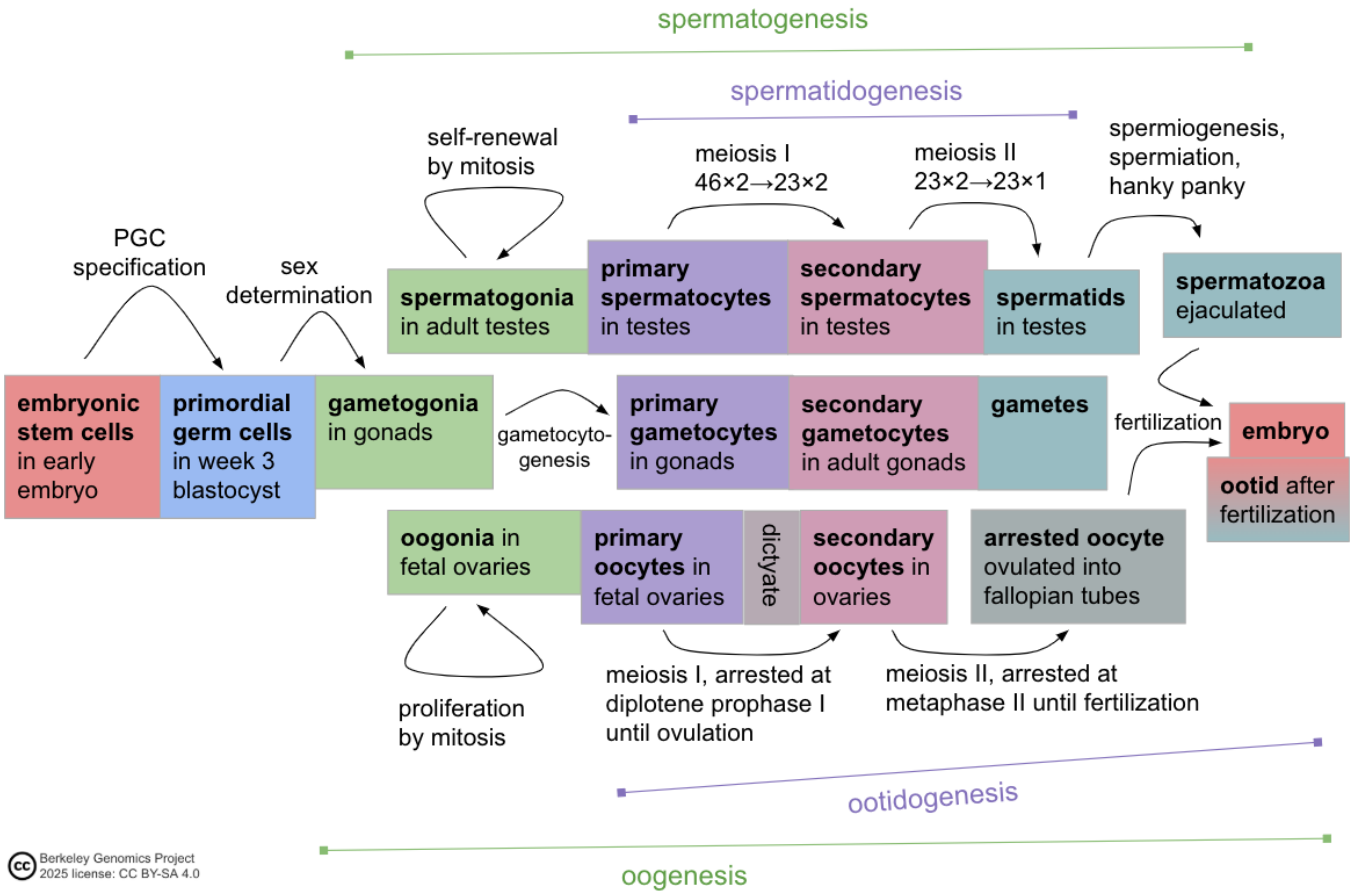


(Graphical abstract from Ishikura et al. (2021).)

Recall the diagram of gametogenesis, with spermatogenesis running along the top half:

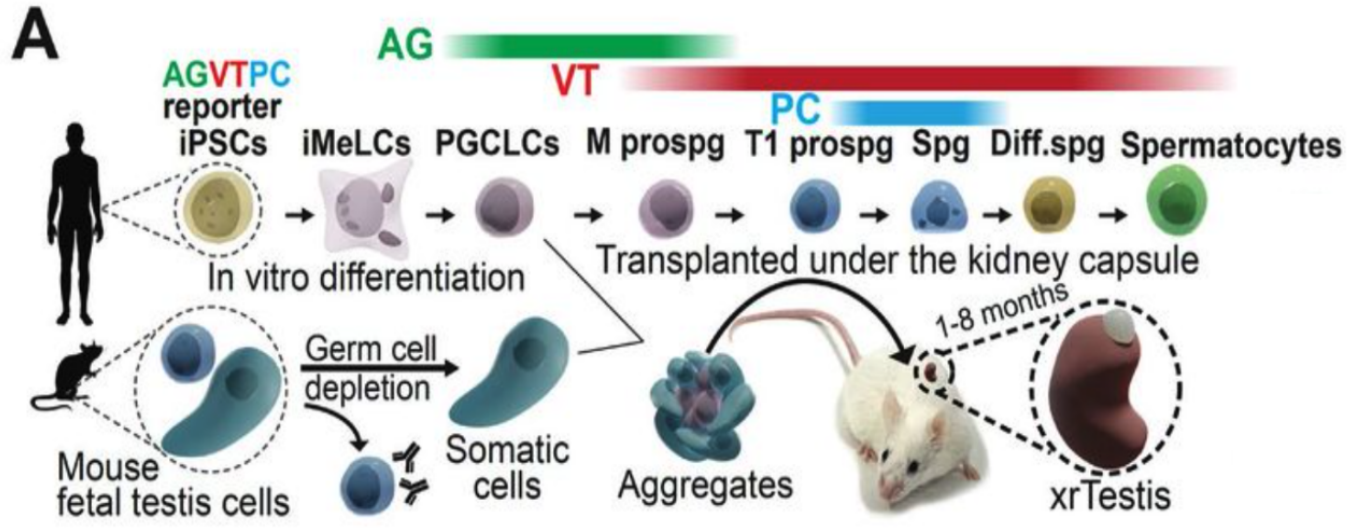
no. 1 (26 July 2023): 12105. https://doi.org/10.1038/s41598-023-39304-1.

[96]Ishikura, Yukiko, Hiroshi Ohta, Takuya Sato, Yusuke Murase, Yukihiro Yabuta, Yoji Kojima, Chika Yamashiro, et al. 'In Vitro Reconstitution of the Whole Male Germ-Cell Development from Mouse Pluripotent Stem Cells'. Cell Stem Cell 28, no. 12 (2 December 2021): 2167-2179.e9. https://doi.org/10.1016/j.stem.2021.08.005.



As we've seen, we can create hiPSCs[97] and thence hPGCLCs[9899]. The next step is to derive spermatogonia.

Recently, Whelan et al. (2024) induced iPSCs to differentiate into spermatogonia-like cells using a somewhat wilder method: they transplanted mouse fetal testicular cells, along with human PGCLCs, into an adult mouse[100]:



(Figure 1A from Whelan et al. (2024).)

The transplanted mouse testicular tissue then recruited blood vessels to feed it, and furthermore *regrew*

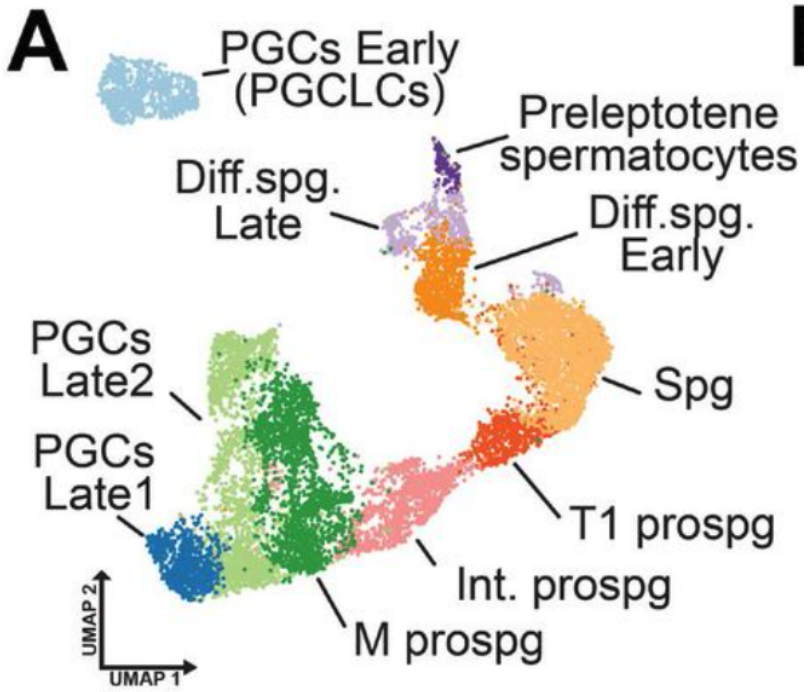
[97]Takahashi, Kazutoshi, and Shinya Yamanaka. 'Induction of Pluripotent Stem Cells from Mouse Embryonic and Adult Fibroblast Cultures by Defined Factors'. Cell 126, no. 4 (25 August 2006): 663–76. https://doi.org/10.1016/j.cell.2006.07.024.

[98]Irie, Naoko, Leehee Weinberger, Walfred W.C. Tang, Toshihiro Kobayashi, Sergey Viukov, Yair S. Manor, Sabine Dietmann, Jacob H. Hanna, and M. Azim Surani. 'SOX17 Is a Critical Specifier of Human Primordial Germ Cell Fate'. Cell 160, no. 1–2 (15 January 2015): 253–68. https://doi.org/10.1016/j.cell.2014.12.013.

[99]Sasaki, Kotaro, Shihori Yokobayashi, Tomonori Nakamura, Ikuhiro Okamoto, Yukihiro Yabuta, Kazuki Kurimoto, Hiroshi Ohta, et al. 'Robust In Vitro Induction of Human Germ Cell Fate from Pluripotent Stem Cells'. Cell Stem Cell 17, no. 2 (6 August 2015): 178–94. https://doi.org/10.1016/j.stem.2015.06.014.

[100]Whelan, Eoin C., Young Sun Hwang, Yasunari Seita, Ryo Yokomizo, N. Adrian Leu, Keren Cheng, and Kotaro Sasaki. 'Generation of Spermatogonia from Pluripotent Stem Cells in Humans and Non-Human Primates'. bioRxiv, 6 May 2024. https://doi.org/10.1101/2024.05.03.592203.

3D tube structure (which I hear is a normal enough result when scientists do this sort of transplant, but still!). Many of the hiPGCLCs differentiated into spermatogonia in the xenografted human-testes-like tissue. Based on scRNA-seq data, a small fraction of the human spermatogonia-like cells appeared to enter the preleptotene substage of prophase I, at the very beginning of meiosis I:



(Figure 2A from Whelan et al. (2024)[101].)

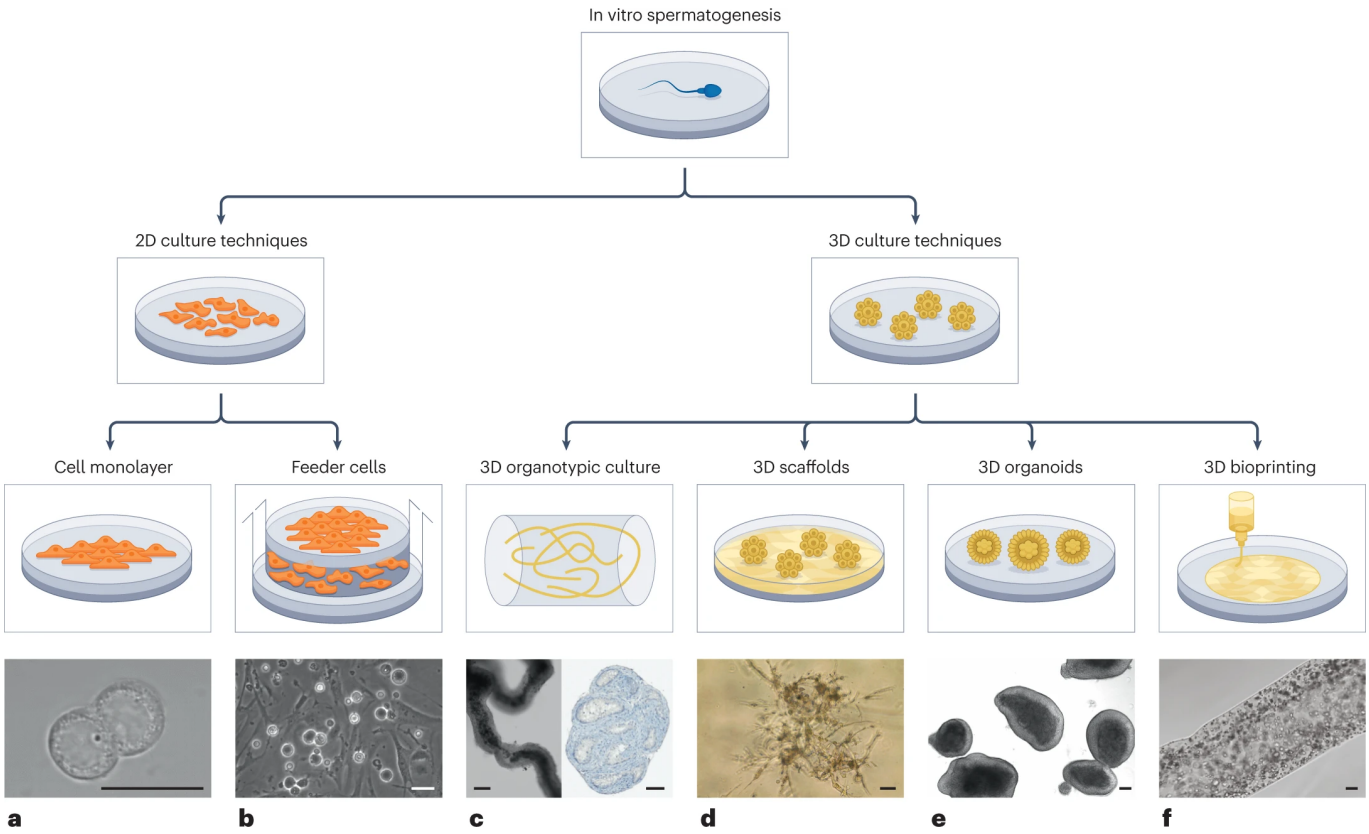
This method is complicated. Also, it seems that their spermatogonia-like cells may not have started to establish paternal imprints. If that's the case, then at best the cells are more like prospermatogonia. Recall that paternal imprinting likely starts to be established already by the time prospermatogonia are differentiating into spermatogonia.

There have been several studies attempting to recapitulate human spermatogenesis in vitro. Specifically, they're attempting to derive from spermatogonia or primary spermatocytes, in vitro, some haploid cells (round spermatids, elongated spermatids, or spermatozoa) that are competent to contribute to an embryo. See the reviews by Tanaka and Watanabe (2023)[102] and Robinson et al. (2023)[103]. An illustration of types of IVS methods that have been tried:

[101]Whelan, Eoin C., Young Sun Hwang, Yasunari Seita, Ryo Yokomizo, N. Adrian Leu, Keren Cheng, and Kotaro Sasaki. 'Generation of Spermatogonia from Pluripotent Stem Cells in Humans and Non-Human Primates'. bioRxiv, 6 May 2024. https://doi.org/10.1101/2024.05.03.592203.

[102]Tanaka, Atsushi, and Seiji Watanabe. 'How to Improve the Clinical Outcome of Round Spermatid Injection (ROSI) into the Oocyte: Correction of Epigenetic Abnormalities'. Reproductive Medicine and Biology 22, no. 1 (9 February 2023): e12503. https://doi.org/10.1002/rmb2.12503.

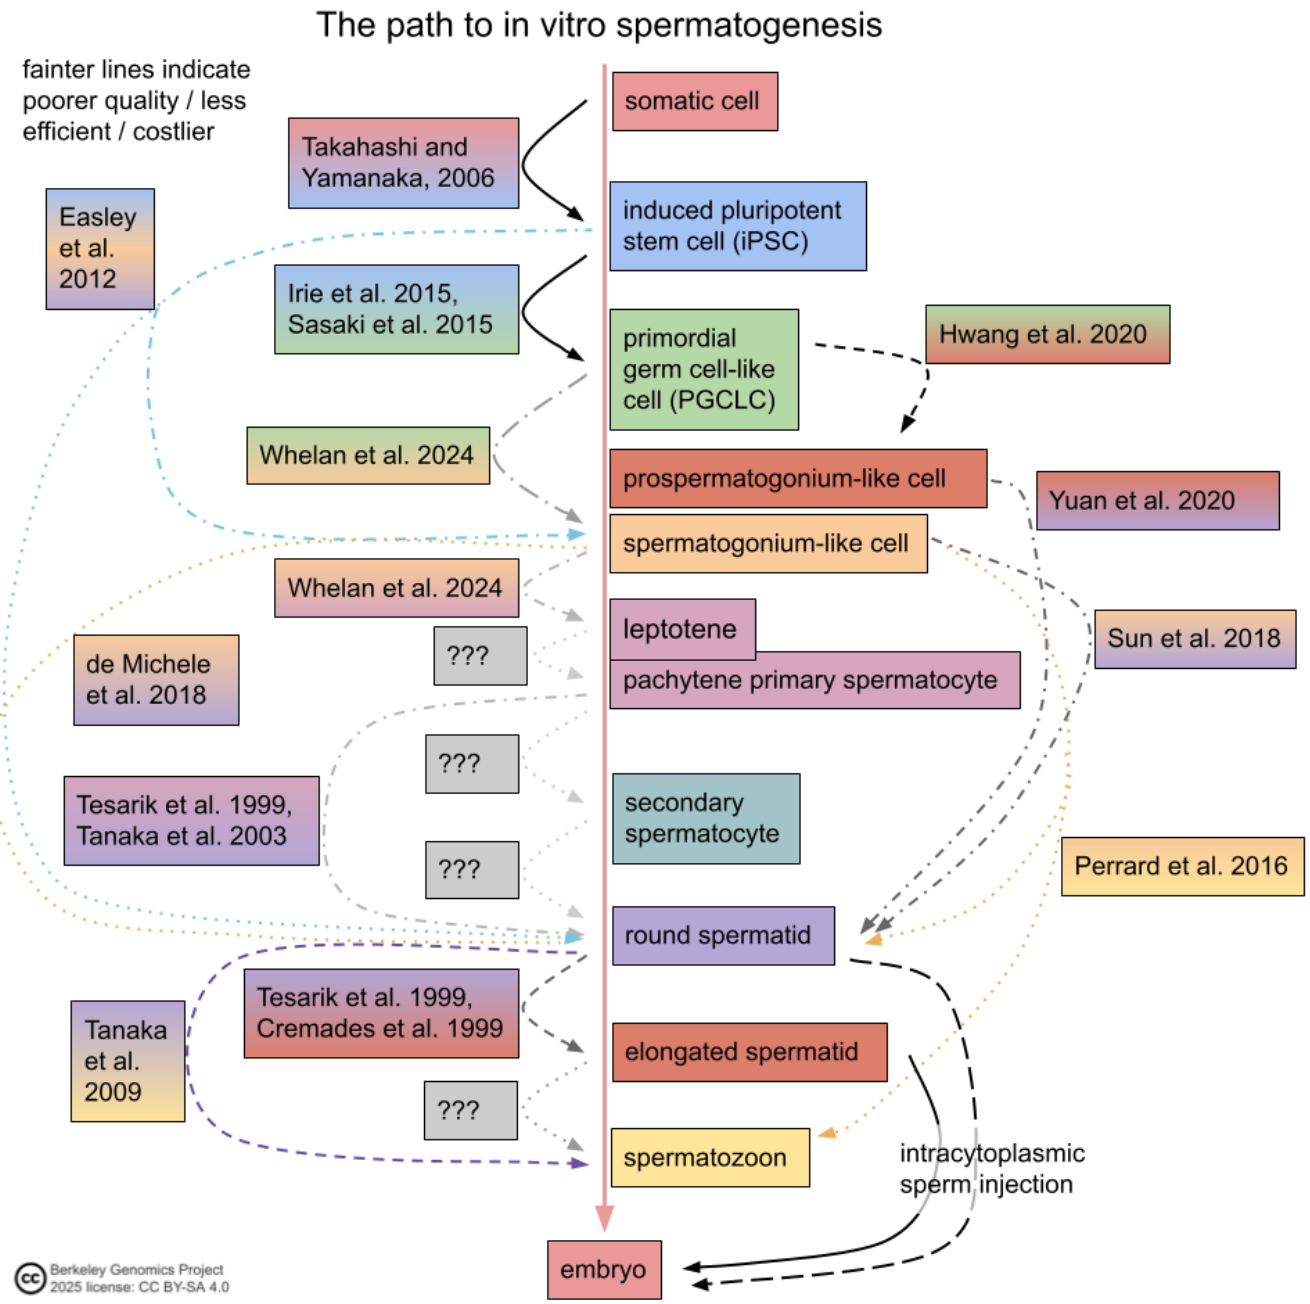
[103]Robinson, Meghan, Sydney Sparanese, Luke Witherspoon, and Ryan Flannigan. 'Human in Vitro Spermatogenesis as a Regenerative Therapy — Where Do We Stand?' Nature Reviews Urology 20, no. 8 (August 2023): 461–79. https://doi.org/10.1038/s41585-023-00723-4.



(Figure 2 from Robinson et al. (2023).)

I summarize some of these studies in the appendix "In vitro spermatogenesis studies".

Below is a diagram summary of the paper summaries, sketching the state of the art of IVS in humans. Down the middle is the progression of in vitro recapitulated male germline development, through spermatogonia, spermatocytes, and then haploids (spermatids, spermatozoa). The arrows show the transitions achieved in vitro, and are labeled with the study. The fainter the arrow, the worse overall utility of the protocol from the perspective of making usable gametes. Most importantly, epigenomic correction for sperm, especially paternal imprinting, has not been achieved with a high degree of correctness.



A potential key issue for in vitro spermatogenesis is that testes have a specific 3D structure that supports spermatogenesis. There's a spatial structure to gene expression within testicular tissue[104], and primary spermatocytes cross a boundary called the Sertoli cell barrier before proceeding through meiosis. This means that, compared to IVO, IVS might require a more complicated method to mimic the 3D structure of testis tissue—hence all the IVS experiments with various kinds of culturing methods.

The paternal EC-making gap remains: we haven't been able to make iPSC cells develop into haploids with correct paternal imprinting.

Using natural epigenomic correction

The idea of natural EC-making is to piggyback off of natural processes that convert non-gamete DNA into gamete DNA. Generally, this would work like this:

1. Genomically vector some type of stem cell in vitro.
2. Transplant the GV cells into gonads or gonad-like tissue.
3. Wait.
4. Retrieve the resulting gametes.

[104]Chen, Haiqi, Evan Murray, Anubhav Sinha, Anisha Laumas, Jilong Li, Daniel Lesman, Xichen Nie, et al. 'Dissecting Mammalian Spermatogenesis Using Spatial Transcriptomics'. Cell Reports 37, no. 5 (2 November 2021). https://doi.org/10.1016/j.celrep.2021.109915.

The *point* of using natural EC-making is that it could be feasible sooner than artificial EC-making. Generally the reason is that, at least in principle, we don't have to know as much before we can do natural EC-making, compared to how much have to know before we can do artificial EC-making. So we wouldn't have to wait for as many large research projects to succeed. There are two reasons:

1. With natural EC-making, we don't have to come up with methods that coax cells through germline development. The natural gonadal tissue already knows how to do that.
2. With natural EC-making, we don't necessarily have to know what constitutes an EC state, any more than we have to verify the epigenomic health of every baby who's conceived today by natural means.

No clear boundary between natural and artificial EC-making

These methods are somewhat contiguous with IVG, in that one kind of approach to IVG aims to induce differentiation, meiosis, and maturation in gametogonia by using one or more of:

- in vivo gonads (i.e. a person's ovaries or testicles),
- gonadal tissue (e.g. extracted from adults, or taken from fetuses who died),
- reconstituted gonads (lumping together unstructured cells from gonads in vitro),

- xenogeneic reconstituted gonads (e.g. using mouse ovarian cells to culture human oogonia),
- transplanted/induced in vivo reconstituted gonadal tissue (e.g. growing human testicular tissue inside a mouse by injecting human stem cells or gonadal cells), or
- gonadal organoids (artificially differentiating stem cells into gonadal cells, e.g. granulosa-like cells).

There's no clear boundary between IVG and natural epigenomic correction, and there's no need for one. But there is a fuzzy spectrum between artificial and natural epigenomic correction, that is useful for orienting in terms of "possibly feasible sooner but slow and costly" vs "feasible later but fast and inexpensive":

EC-making = making cells epigenomically correct for reproduction	
IVG / fully artificial EC-making	**in vivo gametogenesis / more natural EC-making**
in vitro	in vivo, xeno, ex vivo
based on iPSCs, induced gonadal support cells (or 3D organoids), artifical gene regulation (individual cytokines, transcription factors)	based on gonads of a living organism, or extracted gonads cultured in vivo, or xenotransplantation (gonad-like structures grown by injecting gonadal cells e.g. into a mouse)
faster experimental inner loop (directly activates cell processes such as meiosis, fails or succeeds fast-ish)	slow experimental inner loop (tied to biotic processes)
potentially faster full process (approaches the speed intrinsically necessary for the target cell processes)	speed limit (can maybe speed up, but tied to the programmed biotic schedule, which is based on the whole organism's development)
less expensive	more expensive (especially human tissue is expensive)
easier to monitor (cells are more accessible, though filtering still needed)	harder to monitor (cells are inside an organism or organ; have to be filtered out of more extraneous cells)
more novelty required (have to figure out cell signaling, reconstitute several heterogeneous steps manually)	less novelty required (piggybacking on nature's work)
relies on understanding epigenomic correctness	in some cases (e.g. transplanting spermatogonial stem cells into in vivo testicles) we can be confident it will produce EC cells without understanding EC
0.01 to 0.99; probably can eventually reach high efficiency, high quality, low cost	0 to 0.01; can more easily get some target cells, but hard to scale up
doesn't provide info about the target / standard	does model the target / standard of natural gametogenesis

{width:

78% }

Xenotransplantation (of human gonadal tissue into a non-human animal), for example, is an in-between case. I'm counting ex vivo (culturing natural gonads in vitro after extraction) as being natural EC-making, even though it's not fully natural.

The timing problem with natural EC-making

It sounds nice in theory, but unfortunately there's a big obstacle. Sex-specific epigenetic imprinting seems to be partially established already by the adult stage. In the male germline, adult spermatogonial stem cells seem to already have some of their needed imprinting [105][106]; and in the female germline, imprinting seems to be established during the slow maturation of oocytes and follicles before puberty[107]. Adult gonads are likely to not be able to support the changes in germline cells that perinatal / prepubertal gonads support.

What this means is that the straightforward way of using natural EC-making doesn't work. The straightforward method would be:

1. Use some GV method on iPSCs to make GV cells.
2. Transplant the GV cells into a living person's gonads (perhaps with additional reprogramming, e.g. induction into a PGCLC, gametogonia-like, or primary oocyte-like state).
3. Wait.
4. Retrieve the resulting gametes.

But since natural EC-making happens during childhood, this can't be used as an EC-making method because it would involve unnecessary, invasive, dangerous experiments on children.

Alternative gonadal tissue

Are there other ways to get gonadal tissue that can support natural-ish EC-making? Some hypothetical possibilities:

- **Adult gonads for full EC-making.**
  - It is possible that adult gonads might support maturation of immature germline cells.
  - In other words, I don't know that it's impossible for oogonia-like cells, transplanted into a woman's ovaries, to mature into germinal vesicle oocytes. Probably not (and someone might know that it's not possible), but maybe. One observation that suggests it might be possible is that ovaries tend to contain oocytes and follicles at range of different growth stages, which weakly suggests that maybe larger steps in oocyte development can be taken in adult ovaries (though it could very well be that in fact all oocytes in adult ovaries already have their imprinting, or the ones that don't cannot gain imprinting in adulthood). If this is possible, it could follow an application of Yamashiro et al. (2020)'s mouse fetal ovary culture method to make oogonia-like cells from hiPSCs[108]. There would still be a major issue of retrieving the genomically vectored oocytes—normally only a few hundred are ovulated in a woman's life.
  - Likewise, I don't know that it's impossible for prospermatogonia to differentiate into adult spermatogonial stem cells in an adult man's testicles. If so, then it could follow some descendant of Hwang et al. (2020)'s protocol for making prospermatogonia-like cells[109]. Some counter-evidence is Sosa et al. (2018)'s experiments. They transplanted rhesus macaque PGCLCs into adult rhesus testicles. They found that the injected cells made some progress differentiating toward prospermatogonia, but lacked at least one marker of prospermatogonia and didn't contribute to

[105]Trasler, Jacquetta M. 'Epigenetics in Spermatogenesis'. Molecular and Cellular Endocrinology, Proceedings of the 15th Workshop on Molecular and Cellular Endocrinology of the Testis, 306, no. 1 (10 July 2009): 33–36. https://doi.org/10.1016/j.mce.2008.12.018.

[106]Kelsey, Gavin, and Robert Feil. 'New Insights into Establishment and Maintenance of DNA Methylation Imprints in Mammals'. Philosophical Transactions of the Royal Society B: Biological Sciences 368, no. 1609 (5 January 2013): 20110336. https://doi.org/10.1098/rstb.2011.0336.

[107]Trasler, Jacquetta M. 'Gamete Imprinting: Setting Epigenetic Patterns for the next Generation'. Reproduction, Fertility and Development 18, no. 2 (2006): 63. https://www.publish.csiro.au/rd/Fulltext/RD05118.

[108]Yamashiro, Chika, Kotaro Sasaki, Shihori Yokobayashi, Yoji Kojima, and Mitinori Saitou. 'Generation of Human Oogonia from Induced Pluripotent Stem Cells in Culture'. Nature Protocols 15, no. 4 (April 2020): 1560–83. https://doi.org/10.1038/s41596-020-0297-5.

[109]Hwang, Young Sun, Shinnosuke Suzuki, Yasunari Seita, Jumpei Ito, Yuka Sakata, Hirofumi Aso, Kei Sato, Brian P. Hermann, and Kotaro Sasaki. 'Reconstitution of Prospermatogonial Specification in Vitro from Human Induced Pluripotent Stem Cells'. Nature Communications 11, no. 1 (9 November 2020): 5656. https://doi.org/10.1038/s41467-020-19350-3.

spermatogenesis[110].

- **Xeno gonads for full EC-making.**
  - A hypothetical possibility is to use non-human perinatal primate gonads, in vivo. In other words: transplant human iPGCLCs or gametogonia-like cells into a baby monkey's gonads. You might have to first kill the germline cells in the monkey's gonads. Then you wait, and later retrieve matured human germline cells. As long as the human germline cells have matured past the point where imprinting has been established, this would suffice as an EC-making method; it could be followed with in vitro maturation of the gametes.
  - This method might take something like a couple years, depending on the species of non-human primate used. E.g. marmosets take around 1.5 years and cynomolgus macaques take 4 or 5 years to sexually mature, which would be the analog of the human imprinting stage.
  - There are still ethical issues with a protocol like this, but they aren't that bad. The monkey could lead a long and healthy life, except for not being able to have genetic children. However, I'm not sure when—prenatally or postnatally—the hPGCLCs would have to be injected. If they'd have to be injected before natural birth, then e.g. a premature caesarean section would hurt the mother and put the baby at risk.
  - Of course, there would be risks that the epigenetic imprinting, and generally the epigenomic state, wouldn't be correct. There's some theoretical reason to think that the imprinting might be ok: probably the imprinting is normally established in detail by intracellular regulatory processes, which only have to be coarsely activated exogenously. In other words, it's likely enough that the support cells surrounding a germline cell are guiding the germline cell through activations of its own machinery, and they don't have to exert detailed control over the germline cell's epigenetic state.
  - That said, this sort of method would have to at least be validated with epigenetic sequencing. So it would require gold standard measurements—epigenetic data about natural human gametes. Thus it can only possibly bypass half of the EC problem (making the right epigenomic state) and can't bypass the other half (knowing what the EC state is).
  - For oogenesis in particular, this method is extra suspect. During oocyte and follicle maturation, oocytes form cytoplasmic bridges with the adjacent cumulus granulosa cells. The granulosa cells support the oocyte with nutrients, and coregulate with the oocyte. It's conceivable that the resulting oocyte would be fine, and we could validate the results at least in terms of on-DNA epigenomics and general RNA and protein content. But it seems at the very least a major cause for concern. In spermatogenesis, by contrast, germline cells are influenced by signals from surrounding gonad cells, and form bridges with each others as they proliferate, but don't form cytoplasmic bridges with non-germline cells.
- **Long-term culture of fetal gonads.**
  - IVG researchers sometimes use gonads extracted from human fetuses that were aborted or miscarried. If it were possible to expand these tissues in vitro and maintain their fetal-gonad-like function indefinitely, they could be scaled up and used for EC.
  - This is kinda sci-fi though. I don't know of tissues being expanded and maintained in vitro for multiple years. (Several months, though, is not unheard of, and immortalized stem cell lines are used. And of course tissues live in organisms for years—though not developmentally frozen.)

A more thorough literature search that what I've done is called for. See the reviews already cited[111][112][113][114],

[110]Sosa, Enrique, Di Chen, Ernesto J. Rojas, Jon D. Hennebold, Karen A. Peters, Zhuang Wu, Truong N. Lam, et al. 'Differentiation of Primate Primordial Germ Cell-like Cells Following Transplantation into the Adult Gonadal Niche'. Nature Communications 9 (17 December 2018): 5339. https://doi.org/10.1038/s41467-018-07740-7.

[111]Saitou, Mitinori, and Hidetaka Miyauchi. 'Gametogenesis from Pluripotent Stem Cells'. Cell Stem Cell 18, no. 6 (June 2016): 721–35. https://doi.org/10.1016/j.stem.2016.05.001.

[112]Saitou, Mitinori, and Katsuhiko Hayashi. 'Mammalian in Vitro Gametogenesis'. Science 374, no. 6563 (October 2021): eaaz6830. https://doi.org/10.1126/science.aaz6830.

[113]Tanaka, Atsushi, and Seiji Watanabe. 'How to Improve the Clinical Outcome of Round Spermatid Injection (ROSI) into the Oocyte: Correction of Epigenetic Abnormalities'. Reproductive Medicine and Biology 22, no. 1 (9 February 2023): e12503. https://doi.org/10.1002/rmb2.12503.

[114]Robinson, Meghan, Sydney Sparanese, Luke Witherspoon, and Ryan Flannigan. 'Human in Vitro Spermatogenesis as a Regenerative Therapy — Where Do We Stand?' Nature Reviews Urology 20, no. 8 (August 2023): 461–79. https://doi.org/10.1038/s41585-023-00723-4.

and other recent reviews such as [115], [116], and [117] (which I haven't read).

Using natural reproductive DNA

The general idea is to use DNA from natural reproductive cells (sperm, egg, zygote, embryonic stem cell). That DNA will be in a correct epigenomic state. One can then apply a GV method that only does operations that don't mess up the epigenomic state, and then use the resulting cell as a gamete or zygote.

These methods are not EC-making methods—they don't correct anything about the epigenome. They solve epigenomic correctness without doing epigenomic correction. One could call them "EC bypass methods".

Examples:

- simple embryo selection;
- gamete selection;
- single-round embryo prime/base editing;
- **chromosome selection on spermatozoon DNA, mature oocyte DNA, or ESC DNA**;
- iterated CRISPR editing on naive ESCs with imprint maintenance;
- iterated CRISPR editing of SSCs in vivo ("magic rainbow sperm");
- chromosome selection on fully grown germinal vesicle primary oocytes, in theory;
- **iterated CRISPR editing or chromosome selection** on spermatogonial stem cells, with testis transplantation;
- chromosome selection on immature oocytes, with ovarian transplantation.

Natural EC interrupt methods

The last two methods listed above are border cases, which could be called "EC interrupt" methods. They operate on SSCs or immature oocytes—cells which have already been partially but not fully epigenomically corrected. They use natural, partially-EC reproductive DNA, and they also later apply natural EC-making / piggybacking methods.

Natural EC interrupt methods "pause" natural gametogenesis to do GV, before resuming natural EC where it was paused. They don't do any EC-making of their own.

Chromosome selection on grown primary oocytes is an extra strange border case, where the cells have already received all or almost all of their maternal imprinting, but have not yet meiosed. So technically the natural EC-making has completed, but the broader gametogenesis process, which includes meiosis, has to be continued.

An EC interrupt method requires methods for:

1. Obtaining the partially EC cells to apply GV to. E.g. gonadal tissue biopsy.
2. Maintaining whatever partial imprints and cell-type niche that the cell already has, while applying GV. See the next subsection.
3. Completing epigenomic correction after GVing. E.g. retransplanting the GVed cells.

Imprint maintenance problems

EC bypass methods, which use natural EC DNA, apply GV to a cell that is already completely corrected or, in interrupt methods, partially corrected. The imprints that are already established have to be maintained during GV. The requirements and solutions for imprinting maintenance presumably depend on context (which sort of imprints you have to maintain, and what else you're doing to the cells).

For example, MacCarthy et al. (2024) found a modification of the SOX Yamanaka factor that makes the Yamanaka-like cocktail induce and maintain a naive ESC state[118]. In theory you could edit naive ESCs over several rounds while maintaining the imprinting, and then just use one of the edited cells as an embryo.

[115] Cheng, Hanhua, Dantong Shang, and Rongjia Zhou. 'Germline Stem Cells in Human'. Signal Transduction and Targeted Therapy 7, no. 1 (2 October 2022): 1–22. https://doi.org/10.1038/s41392-022-01197-3.

[116] Ibtisham, Fahar, and Ali Honaramooz. 'Spermatogonial Stem Cells for In Vitro Spermatogenesis and In Vivo Restoration of Fertility'. Cells 9, no. 3 (March 2020): 745. https://doi.org/10.3390/cells9030745.

[117] Kulibin, A. Yu, and E. A. Malolina. 'In Vitro Spermatogenesis: In Search of Fully Defined Conditions'. Frontiers in Cell and Developmental Biology 11 (24 February 2023). https://doi.org/10.3389/fcell.2023.1106111.

[118] MacCarthy, Caitlin M., Guangming Wu, Vikas Malik, Yotam Menuchin-Lasowski, Taras Velychko, Gal Keshet, Rui Fan, et al. 'Highly Cooperative Chimeric Super-SOX Induces Naive Pluripotency across Species'. Cell Stem Cell 31, no. 1 (4 January 2024): 127-147.e9. https://doi.org/10.1016/j.stem.2023.11.010.

However, there would be some amount of loss of imprinting due to being cultured and due to the action of super-SOX. (One might be able to then fix the hopefully not-too-numerous aberrations with epigenetic editing.) Further, there's some indication that the rapid reprogramming (e.g. genome-wide demethylation) that occurs in very early embryonic development might already be enough to disrupt an ESC's ability to act as a zygote—e.g. Mitalipov et al. (2002)'s rhesus macaque blastomere SCNT experiments had poor pregnancy rates and one stillbirth (though not necessarily due to developmental abnormality)[119]. On the other hand, Meng et al. (1997) also did SCNT to rhesus macaque blastomeres, and out of 29 embryos in 9 females got one 30-day miscarriage and two apparently healthy live births[120]. More study needed.

Broadly speaking, the key question is: How, and how well, can the epigenomic state be maintained during culture in vitro, while applying some genomic vectoring method? If a near-perfect maintenance method is found, further epigenetic correction wouldn't be needed.

Bayerl et al. (2021) found a culture medium that seemed to maintain imprinting in ESCs cultured for several cell passages, though there was still loss of imprinting within a couple weeks (see figure S3H in the supplemental information)[121]. Tan et al. (2020) were able to characterize and culture SSCs as distinct from differentiating spermatogonia, which hadn't been well-separated before[122]. They didn't report on imprinting status of their cultured SSCs, so we don't know their quality. It's plausible that a "close enough" SSC niche would automatically maintain any important imprinting, but it's also plausible that it wouldn't.

It might be fairly feasible to maintain imprinting in culture. There are natural mechanisms that maintain imprinting during very early embryonic demethylation. The paternal genome is very quickly mostly demethylated, even before the first cell division, but the imprints remain (as well as some other methylations, e.g. to suppress transposons). Some proteins (Stella, Zfp57, and Trim28) are known to be involved in protecting paternal imprints[123]. Fischer et al. (2025) used Zfp57 to protect imprints somewhat during naive hPSC induction, and they discuss other protective proteins[124]. (Though, they also tried Stella (DPPA3) and found no effect—perhaps Tet3 wasn't around?)

Epigenetic CRISPR editing

CRISPR-Cas9 was a breakthrough in DNA editing in part because, in contrast to previous DNA editors with protein-based binding, it could be fairly easily, flexibly, and precisely targeted to approximately any specific region of DNA using synthesized guide RNAs that bind to complementary DNA single strands. A further innovation was to deactivate the elements of Cas9 that cut DNA, creating dCas9, which doesn't cut DNA. You can attach thingies to dCas9 that have various effects on and around the DNA that the system binds to. Some of those thingies will cause some epigenetic state to change (perhaps by recruiting cellular machinery)—e.g. methylate some CpG islands, chemically modify some histones, or affect chromatin accessibility. Various such systems have been developed in the past decade[125][126].

The idea for EC-making with epigenetic CRISPR editing would be to make a GV stem cell, and then apply epigenetic CRISPR to the cell. This could be haploid or diploid:

[119]Mitalipov, Shoukhrat M., Richard R. Yeoman, Kevin D. Nusser, and Don P. Wolf. 'Rhesus Monkey Embryos Produced by Nuclear Transfer from Embryonic Blastomeres or Somatic Cells1'. Biology of Reproduction 66, no. 5 (1 May 2002): 1367–73. https://doi.org/10.1095/biolreprod66.5.1367.

[120]Meng, Li, John J. Ely, Richard L. Stouffer, and Don P. Wolf. 'Rhesus Monkeys Produced by Nuclear Transfer1'. Biology of Reproduction 57, no. 2 (1 August 1997): 454–59. https://doi.org/10.1095/biolreprod57.2.454.

[121]Bayerl, Jonathan, Muneef Ayyash, Tom Shani, Yair Shlomo Manor, Ohad Gafni, Rada Massarwa, Yael Kalma, et al. 'Principles of Signaling Pathway Modulation for Enhancing Human Naive Pluripotency Induction'. Cell Stem Cell 28, no. 9 (2 September 2021): 1549-1565.e12. https://doi.org/10.1016/j.stem.2021.04.001.

[122]Tan, Kun, Hye-Won Song, Merlin Thompson, Sarah Munyoki, Meena Sukhwani, Tung-Chin Hsieh, Kyle E. Orwig, and Miles F. Wilkinson. 'Transcriptome Profiling Reveals Signaling Conditions Dictating Human Spermatogonia Fate in Vitro'. Proceedings of the National Academy of Sciences 117, no. 30 (28 July 2020): 17832–41. https://doi.org/10.1073/pnas.2000362117.

[123]Champroux, Alexandre, Julie Cocquet, Joëlle Henry-Berger, Joël R. Drevet, and Ayhan Kocer. 'A Decade of Exploring the Mammalian Sperm Epigenome: Paternal Epigenetic and Transgenerational Inheritance'. Frontiers in Cell and Developmental Biology 6 (15 May 2018). https://doi.org/10.3389/fcell.2018.00050.

[124]Fischer, Laura A., Brittany Meyer, Monica Reyes, Joseph E. Zemke, Jessica K. Harrison, Kyoung-mi Park, Ting Wang, Harald Jüppner, Sabine Dietmann, and Thorold W. Theunissen. 'Tracking and Mitigating Imprint Erasure during Induction of Naive Human Pluripotency at Single-Cell Resolution'. Stem Cell Reports 0, no. 0 (13 February 2025). https://doi.org/10.1016/j.stemcr.2025.102419.

[125]Nakamura, Muneaki, Yuchen Gao, Antonia A. Dominguez, and Lei S. Qi. 'CRISPR Technologies for Precise Epigenome Editing'. Nature Cell Biology 23, no. 1 (January 2021): 11–22. https://doi.org/10.1038/s41556-020-00620-7.

[126]McCutcheon, Sean R., Dahlia Rohm, Nahid Iglesias, and Charles A. Gersbach. 'Epigenome Editing Technologies for Discovery and Medicine'. Nature Biotechnology 42, no. 8 (August 2024): 1199–1217. https://doi.org/10.1038/s41587-024-02320-1.

- Haploid epigenetic editing: add either maternal imprinting or paternal imprinting, and then use the result as the maternal / paternal DNA contribution to a zygote.
- Diploid epigenetic editing: add both maternal and paternal imprinting to a diploid cell, and then use the result as a zygote or as the DNA contribution to a zygote (via SCNT).
  - However, it might be infeasible to usefully edit a diploid genome this way. By default, at least for autosomes, you'd get edits on both homologous chromosomes, so you wouldn't have the appropriate differentially methylated regions or other monoallelic state.
  - If there happen to be DNA differences between the homologs, e.g. a nearby SNP, those could be used to target just one allele on a specific chromosome.
  - You could calibrate the efficiency of the editor so that on average some fairly small portion of target loci get the epigenetic edit. Then you could do, essentially, iterated epigenetic CRISPR editing—though it seems at least fairly likely that the epigenomic disruption from cell culturing would be worse than the progress you make with editing. This would probably work for only at most a small number of edits, but could be useful in hybrid EC-making.
  - Even if you can get monoallelic edits, in theory you could still have problems if you aren't targeting a specific chromosome. Suppose you make the right epigenetic edit on exactly one chromosome. Plausibly you'd thus get your desired trans-regulatory effects (i.e. long-range regulation, e.g. one gene makes a transcription factor that diffuses around and binds and regulates some other gene on some other chromosome). But you might not get your desired cis-regulatory effects (i.e. effects on nearby DNA, e.g. you methylate a promoter and thereby suppress expression of the gene next to that promoter). Plausibly there wouldn't be any missing effects, e.g. because the DNA would be mostly identically regulated except at a few places, none of which have cis-regulatory interactions with each other.
- Diploid-to-haploid epigenetic editing: You establish biallelic maternal imprinting or biallelic paternal imprinting. Then you somehow derive a haploid without disrupting the state.
  - E.g. possibly you could induce meiosis, though this would probably disrupt the state.
  - E.g. possibly you could do chromosome selection of some sort.
  - I'm not sure if there's ever a benefit to this, over just haploid editing. One possible benefit is that the diploids might be easier to culture than the haploids, making iterated epigenetic CRISPR editing easier.

Challenges:

- CRISPR-based epigenetic editors have many of the same problems as CRISPR-based DNA editors. They generally will have off-target effects—e.g. they methylate regions that weren't targeted[127]. Depending on the variant, they might have low efficiency (they don't make the target edit). They also might be difficult to deliver to the target cells if the cells are embedded in some tissue (e.g. germline-cell-like cells in gonadal tissue).

- Epigenetic editors whose function relies on cellular machinery wouldn't work if the cell being edited doesn't have that machinery. E.g., sperm-like cells would tend to be missing a lot of normal machinery because they are mostly inactive cells.
- Full EC-making would be difficult using epigenetic CRISPR editing alone.
  - In other words, it would probably be hard to take an hiPSC (haploid or diploid), apply some epigenetic editors, and get a cell that can be used as a zygote nucleus or as a uniparental contribution.
  - Full EC-making would require not only establishing sex-linked imprints, but also general epigenomic reprogramming. E.g., PGCs become broadly demethylated compared to normal ESCs, and gametes become very silenced. This would have to be accomplished some way other than a huge number of single locus epigenetic edits.
  - Safe full artificial EC-making, in contrast to natural EC-making, would require knowing in great detail the necessary epigenetic state.
  - With EC-making that is natural or natural-like (e.g. induced in vitro), the imprints are established via something like the natural cell pathways. Those pathways would tend to do all or almost all of the reprogramming needed, somewhat by default, even if you don't know about some of the needed imprints. In contrast, with epigenetic editing, the most upstream epigenetic imprints are each individually likely to be absent unless you specifically make those edits.

[127] Liesenfelder, Sven, Mohamed H. Elsafi Mabrouk, Jessica Iliescu, Monica Varona Baranda, Athanasia Mizi, Martina Wessiepe, Argyris Papantonis, and Wolfgang Wagner. 'Epigenetic Editing at Individual Age-Associated CpGs Affects the Genome-Wide Epigenetic Aging Landscape'. bioRxiv, 5 June 2024. https://doi.org/10.1101/2024.06.04.597161.

- If the goal is to make a haploid genome, the cell would have to be meiosed somehow. This would separately require a method to induce meiosis. If you did the EC-making before meiosis, meiosis might mess up the epigenetic state.

So, epigenetic editing seems more likely to be used in hybrid methods than as a full EC-making method. It could be useful to e.g. fix a few imprints that were lost while culturing naive ESCs or that failed to be established by an imperfect IVG method.

As an illustration of the principle, though not directly an example of epigenetic CRISPR editing: Li et al. (2025) made bi-paternal mice by making genetic edits, targeted to genes affected by sex-linked imprinting, aiming to correct the activity of those genes. This was partly successful, in that they greatly decreased the rates of various abnormalities, and produced live births; but only a few made it to adulthood[128].

A more direct illustration, mentioned above: Liu et al. (2018) cloned cynomolgus macaques with SCNT[129]. They altered the histone modifications in the DNA before transferring it to enucleated oocytes. This greatly improved their SCNT success rates (though they still had epigenetic near-misses, as mentioned above).

Donor embryo

This is a method for lowering the required EC bar. You don't ask for a cell that can form the entire conceptus, including the fetus and placenta and amniotic sac, as a normal zygote can. Instead, you take a donor embryo and somehow disable it from growing a fetus. Then you inject your GVed cell or embryo into the donor conceptus. The injected cells take the fetal role, while the donor grows the placenta and amniotic sac. See for example the VelociMouse method[130].

The main benefit of this method is not exactly about imprinting, but rather about the pluripotency of your GV cells. Naive ESCs can form the whole conceptus, but primed ESCs can only form the fetus and not the trophectoderm. Naivety is harder to maintain in culture than a more differentiated state. By using a donor embryo to provide the trophectoderm-descended tissues, you don't need to maintain naive pluripotency.

Another part of the hope with a donor embryo is that most sex-linked imprinting is about the placenta, rather than the fetus. So, even if the progenitor still requires some EC-making, the amount of EC-making you still have to do might be greatly decreased. For example, it might require fewer epigenetic edits.

I don't know much about this method. I suspect it's not straightforwardly viable because sex-linked imprinting is not just about the trophectoderm. It also might produce mosaicism. The concern about mosaicism might be largely avoided by using a twinned embryo as the donor.

Hybrid natural and artificial EC-making

It might be that we find a method to partially but not fully correct the imprints of some DNA artificially in vitro. In that case, we could compose that method with natural epigenomic correction.

Examples:

- *Natural EC DNA and epigenetic editing.* EC bypass methods that want to just use naturally EC DNA don't necessarily fully preserve imprinting. E.g. culturing and editing ESCs or SSCs would, by strong default, lead to some loss of imprinting. But if they mostly preserve imprinting, the remaining aberrations could maybe be corrected with a small number of epigenetic edits.
- *Spermatogonial stem cell induction.* If we could artificially make iSSCs, we could apply GV to iPSCs, induce iSSCs, and then transplant them. This is a genuinely hybrid EC-making method because SSCs already have some paternal imprinting; our induction method would have to make those imprints, as well as general epigenome-wide reprogramming.
- *In vitro spermatogenesis proper.* Conversely, suppose we find a method to take natural SSCs and induce them into spermatogenesis in vitro, i.e. meiosis and further epigenomic reprogramming. In this

[128] Li, Zhi-kun, Li-bin Wang, Le-yun Wang, Xue-han Sun, Ze-hui Ren, Si-nan Ma, Yu-long Zhao, et al. 'Adult Bi-Paternal Offspring Generated through Direct Modification of Imprinted Genes in Mammals'. Cell Stem Cell, 28 January 2025. https://doi.org/10.1016/j.stem.2025.01.005.

[129] Liu, Zhen, Yijun Cai, Yan Wang, Yanhong Nie, Chenchen Zhang, Yuting Xu, Xiaotong Zhang, et al. 'Cloning of Macaque Monkeys by Somatic Cell Nuclear Transfer'. Cell 172, no. 4 (8 February 2018): 881-887.e7. https://doi.org/10.1016/j.cell.2018.01.020.

[130] Dechiara, Thomas M., William T. Poueymirou, Wojtek Auerbach, David Frendewey, George D. Yancopoulos, and David M. Valenzuela. 'VelociMouse: Fully ES Cell-Derived F0-Generation Mice Obtained from the Injection of ES Cells into Eight-Cell-Stage Embryos'. Methods in Molecular Biology (Clifton, N.J.) 530 (2009): 311–24. https://doi.org/10.1007/978-1-59745-471-1_16.

case, we could extract adult SSCs, apply a GV method, and then use IVS. The GV method would have to preserve the SSC imprinting that is supposed to be transmitted to male gametes.

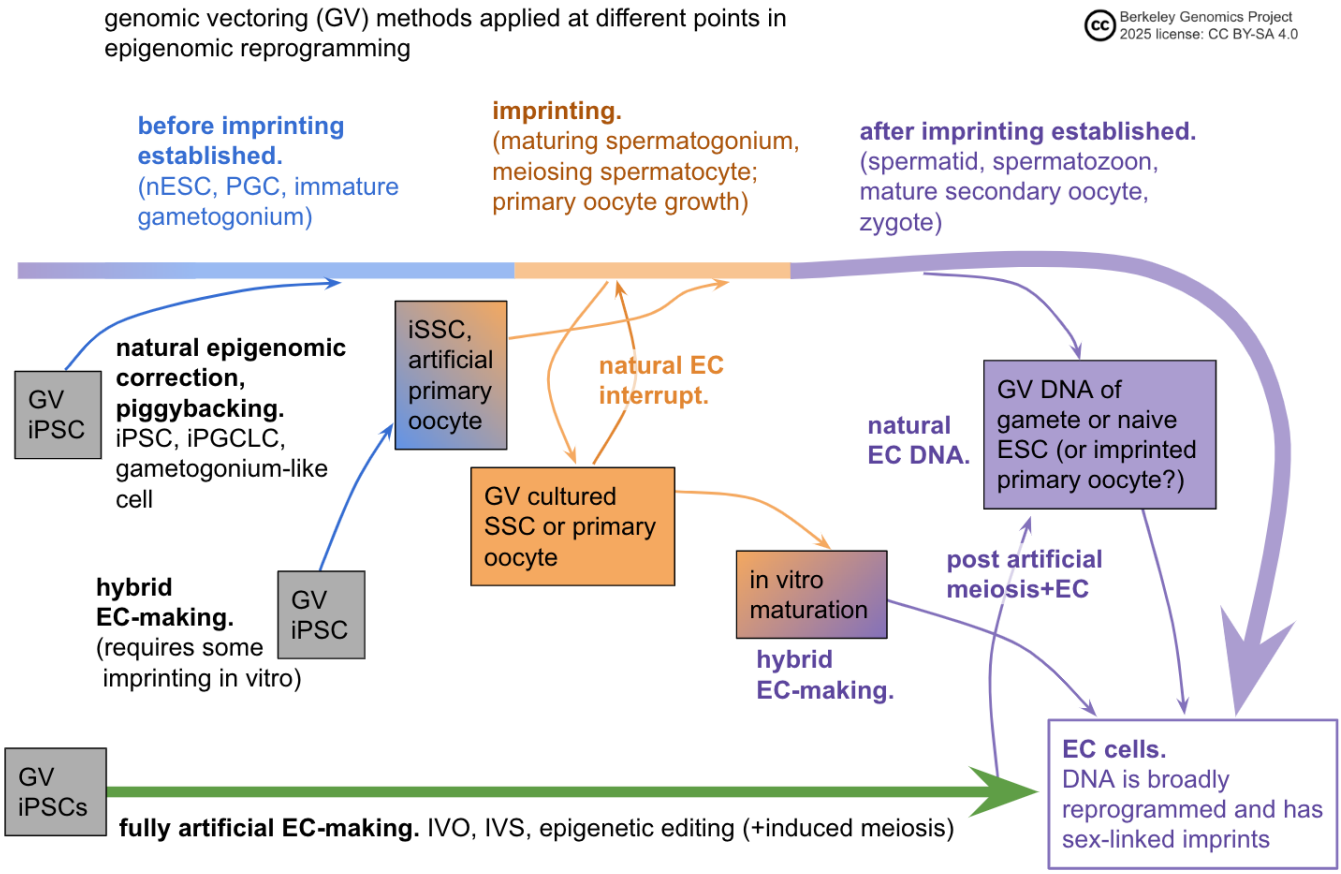
- *Immature oocytes.*
  - In theory, it might be possible to retrieve oocytes from ovaries that are not yet at the germinal vesicle stage with preantral / antral follicles. We could apply a GV method to these, and then apply in vitro maturation to complete the oocyte growth and imprinting, and undergo meiosis.
  - This is probably infeasible, as there's no way to get such immature oocytes in general.
  - At first this hypothetical method seems entirely pointless. E.g., it doesn't seem especially easier to apply GV to immature rather than mature oocytes. But one potential benefit is that, if there were a method to safely and ethically obtain immature oocytes (e.g. if there are some in adult ovaries), then they might be more abundant than mature oocytes. All oocyte maturation methods that I'm aware of have low efficiency, in the sense that there are many more immature than mature ones. So if the GV method has some attrition, it might be better to apply GV at a stage where you get more chances. (However, I don't see a benefit to using pre-growth oocytes rather than just using the kind of immature oocytes that are abundant in adult ovaries, which have already received all or almost all of their imprints.)

How GV and EC interact

Reproductive genomic vectoring requires a GV method and a method for handling epigenomic correctness. Is that sufficient? If we have a viable GV method and EC method, do we get reproductive GV? Not necessarily.

A simple-ish and wrong model of combining GV and EC

Here's a diagram displaying, coarsely, ways that GV can interface with EC:



Basically, at a coarse level, you can apply GV at some point in the epigenomic correction process. This could be before imprinting is established; while it's being established; or after it's been established.

GV and EC are not conceptually separable

It would make it easier to think about reproductive GV if we could think about GV and EC separately. Then we could solve GV with some method, and separately solve EC with some other method, and then

combine them to get a cell that has been genomically vectored as we choose and that is epigenomically competent for development.

To some extent, this can be done. For example, suppose we had full IVO: we can take any diploid stem cell and make a competent egg. Now, it seems like basically any GV method whatsoever can be precomposed with our IVO method. We take a somatic cell from the parent, induce it to an iPSC, apply our GV method to it, and then use IVO on the GVed iPSC to get a GVed egg.

But there is a caveat missing. It's *not* technically true that any GV method could be used in exactly this way. Specifically, consider recombinant chromosome selection. The power of recombinant chromosome selection is that recombined chromosomes can be higher scoring than either of the two parent chromosomes. So it only makes sense to do recombinant chromosome selection *after meiosis*. But taken literally the above protocol would say to "do recombinant chromosome selection on iPSCs, then do IVO", which doesn't make sense, as iPSCs are diploid and are genomically the same as the adult whose cells they're derived from.

This isn't some great difficulty with recombinant chromosome selection. You could probably, for example, do chromosome selection to some artificial eggs to make a diploid; make it an iPSC; and do IVO again to get a high-scoring, EC oocyte.

My point is rather that, in general, we can't just separate out GV from EC.

More examples to illustrate the point:

- Chromosome selection might be doable on sperm DNA. However, iterated editing almost certainly can't be done on sperm DNA: sperm don't have the requisite cell machinery, sperm DNA is highly inaccessible, and sperm can't be cultured for iteration. Thus, the general EC bypass method "just use natural sperm DNA" is only compatible with some GV methods. Likewise for natural oocytes, as oocytes have already started meiosis and therefore can't normally be cultured.
- On the other hand, if naive ESCs can truly be cultured such that they would still form viable embryos, then both editing and chromosome selection ought to work with naive ESCs. So *some* EC bypass methods *could* work with both editing and selection methods.
- Suppose you can culture adult spermatogonial stem cells in a way that maintains their imprinting. Then, you can do iterated editing to these SSCs, transplant them back into the adult's testicles, and retrieve edited spermatozoa. Can you also do chromosome selection? On the face of it, yes; chromosome selection is supposed to be minimally disruptive of epigenetics, in theory. However, chromosome selection requires access to individual cells, separated from their context. If the epigenomic state of SSCs is fragile, this separation would be disruptive. So chromosome selection may not be compatible with this EC method. Then again, editing also wants to separate the SSCs for a little while in order to verify edits and non-damage; but then again again, maybe that step can be deferred to a single screening round of embryos after retrieval and fertilization.
- Different implementations of broad methods have different epigenomic consequences. For example, chromosome selection that uses some kind of mechanical manipulation might be fast enough to not damage epigenomic state. On the other hand, poor man's chromosome selection involves several rounds of mitosis and exotic tetraploid states, so it seems unlikely to preserve epigenomic state by default. Both would work on iPSCs followed by IVO; but only mechanical chromosome selection seems plausibly compatible with using natural SSCs.

Features of GV and EC methods that affect compatibility

A given GV or EC method is, vaguely speaking, some strategy for applying some manipulations to some cells or DNA. Here are some aspects of GV or EC manipulations that affect which other methods they can and can't be combined with:

- Ploidy.
  - Some methods have to work with haploid cells, and some have to work with diploid cells.
  - IVG methods probably can't apply straightforwardly to haploid cells.
  - E.g. if poor man's chromosome selection works via spontaneous tetraploid reduction, then it has to work with diploids.
  - E.g. epigenetic CRISPR editing to make a gamete has to work with haploids (unless imprints can be preserved across meiosis).
  - Methods can preserve (e.g. mitosis) or change (e.g. meiosis, cell fusion) ploidy. Most IVG protocols probably can't perform imprinting without also turning diploids into haploids.

  - Meiosis is not necessarily the only way to make haploid cells. For example, some forms of chromosome selection might allow one to make a haploid "manually". E.g. if you could specifically destroy several target chromosomes without totally destroying a cell, you could destroy a haploid genome, leaving the other haploid genome.
  - GV methods applied before diploid-haploid reduction will generally be half as powerful (if the haploid is used).
- Position relative to meiosis.
  - By default, EC-making methods have to either come after meiosis, or come along with meiosis (most IVG), or have a way of preserving imprints across artificial meiosis.
  - Recombinant chromosome selection has to come after at least one step of meiosis.
- Culturability.
  - Some methods can only work with culturable cells.
  - Generally iterated methods, such as iterated editing, iterated meiotic selection, and poor man's chromosome selection, only work with cells that can be grown.
  - Mechanical chromosome selection could work with non-culturable cells.
  - Some methods produce non-culturable cells.
  - Generally, meiosis produces haploid cells, which might not be easily culturable.
  - EC-making methods might produce cells that can't be cultured by default, but rather can only be cultured if there's an additional method to maintain imprints.
  - Fully EC cells are heavily silenced, and therefore can't be cultured.
  - Most IVG methods produce cells with naturally very silenced genomes, which therefore can't be cultured without further intervention.
  - Some selection methods might produce intermediate diploid cells with a lot of homozygosity, which might be unable to mitose or might be unsafe to mitose (because you're applying some weird selection pressure).
- Some methods don't automatically preserve epigenomic correctness.
- Cell access.
  - Some methods require total access to isolated cells in vitro, some require less total access.
  - Editing doesn't require total access. E.g. one could edit SSCs that are growing inside testicular tissue in vitro or even in vivo. (However, without at least one step of total access for screening, some DNA damage accumulates; so you have to screen embryos, which is costly.)
  - Most selection methods require total access.
  - Some methods provide access, some don't.
- Cell stress.
  - Cells get stressed by damage. Then they aren't as good at proliferating or other activities.
  - Most GV methods would stress the GVed cell.
  - EC methods can be more or less compatible with stressed cells.
  - For example, I would guess that sperm being stressed is not so bad. Sperm are mostly inert anyway, and are dissolved upon entering the oocyte, so they don't have much responsibility to add to their stress lol. Therefore, e.g. sperm selection (using natural EC DNA as an EC bypass method) should not be a problem in terms of stress.
  - On the other hand, applying a GV method to SSCs and then retransplanting them into testes would pit the GVed SSCs against the natural SSCs already in the testes. If the GVed SSCs are stressed, then by default they would likely be outcompeted by the other SSCs.
- Some methods such as genetic or epigenetic CRISPR editing require some machinery to be present in the cell, e.g. methylation enzymes.
- Some methods require certain chromatin states.
  - Chromosome selection requires that the chromosomes be condensed enough to manipulate, but not entangled with or bound to each other (as in meiosis).
  - CRISPR editing and DNA labeling require that the chromatin be accessible.

How feasible are different epigenomic correctness methods?

Most of the following judgements are speculations. They're based on reading what biologists write, but many of them are just guesses; they're intended as a snapshot of my current understanding, and to give experts something to disagree with.

These are arranged in rough order of easiest to hardest, in terms of getting them to work as methods of handling EC (but *not* necessarily as strong GV methods).

- Just use normal reproductive DNA.
  - E.g. simple embryo selection, sperm or oocyte chromosome selection.
- Just use normal reproductive DNA, with a short window of operating on the DNA.
  - E.g. a small amount of editing on ESCs; ESC chromosome selection.
- Use natural epigenomic correction.
  - E.g. in vivo spermatogonial editing; spermatogonial stem cell induction and testicle transplantation; oogonial stem cell induction followed by oocytogenesis and maturation in ovarian tissue in vitro.
- Imprint maintenance and normal reproductive DNA, plus a small amount of epigenetic editing.
  - E.g. editing ESCs for a longer period of time.
- IVO.
- IVS.
- Donor embryo. (I'm guessing donor embryo methods are easier to make sort-of work, but also more likely to be a near-miss.)
- Massive epigenetic editing to directly create a gamete or embryo.
- Creating an embryo directly from a stem cell in some other way.

How epigenomically disruptive are different GV methods?

Note: **The following judgements are speculations**, like the ones above about feasibility of EC methods.

No epigenomic disruption

These methods might not disrupt epigenomic state at all:

- Simple cell population selection
- Mechanical chromosome selection
- Single-round multiplex CRISPR editing

Thus, these methods can use EC bypass methods, e.g. simple embryo selection (using already EC cells), editing zygotes, or sperm chromosome selection. These methods are most likely to be feasible now or in the near future.

Little epigenomic disruption

These methods might not disrupt epigenomic state too much:

- Iterated multiplex CRISPR editing
  - This method requires several weeks or months to operate, and requires cells to undergo mitosis. During that time, by default, special epigenomic states such as being a naive ESC or being a spermatogonium will degrade, e.g. with loss of sex-linked imprinting or loss of stem cell potency.

- Chromosome selection via chromosome digestion and electrofusion
  - This theoretical method might totally destroy cells or massively disrupt their state by setting off lots of DNA damage response mechanisms. But if that can be circumvented somehow, the intact chromosomes might be mostly unaffected.

So these methods can take advantage of weak EC-making methods, e.g. small numbers of epigenetic CRISPR edits or transplanting vectored gametogonia back into in vivo gonads.

Significant epigenomic disruption

These methods probably disrupt epigenomic state a lot:

- Very many rounds of iterated multiplex CRISPR editing
- Iterated recombinant selection
- Chromosome selection via whole cell fusion and tetraploidy reduction

Many rounds of CRISPR editing involve even more cell culturing, which by default disrupts epigenomic state. For example, if culturing takes a couple weeks, and you can get 20 edits per round, and you want several hundred edits, that takes over half a year. That said, it may be feasible to find methods, in the vein of super-SOX, to strongly maintain the needed epigenomic state in culture, in which case strong iterated editing would work.

Iterated selection, e.g. IMS, involves inducing meiosis. That process presumably has significant epigenomic effects, including some loss of imprinting. More importantly, because IMS shuffles chromosomes and segments of chromosomes between cells, it will mix and match chromosome segments with paternal and maternal epigenetic imprinting. This means that IMS absolutely needs a strong EC-making method.

Tetraploidy may not cause further significant epigenomic disruption even with methods to maintain imprinting, so it may require strong EC-making methods.

Haploid vs. diploid

Some ways of handling epigenomic correctness allow for a diploid GV genome to contribute to an embryo; others only allow for a haploid contribution.

In general, a haploid GV genome contribution has half the GV power as a diploid contribution:

- For editing, most edits will be fairly common variants (minor allele frequency >1%, say). For such edits, in most cases it's fine to edit both chromosomes. Homozygosity shouldn't be a problem for common variants. So usually a diploid contribution can take twice as many edits as a haploid one, giving twice the effect, for roughly the same cost (most of the cost is one-time-per-edit-type, e.g. figuring out a good guide RNA to use).
- For selection, the effects simply add between the two halves of the genome. (That is, selecting half a genome is half the effect of separately selecting both halves. However, selecting half a genome is not half the effect of selecting a whole genome, but rather is $1/\sqrt{2}$ times the effect, i.e. roughly 70%.)

Which ways of handling the EC problem enable haploid vs. diploid GV?

- In vitro gametogenesis is a haploid EC-making method. Just because you can do in vitro oogenesis, doesn't necessarily mean you can do in vitro spermatogenesis, and vice versa. So you only get haploid GV power. (If you have both IVO and IVS, then you probably always get diploid GV power.)
- Natural EC-making likewise gives haploid GV power.
- Epigenetic CRISPR editing ought to be diploid, though I don't know. If you have epigenetic editing methods working well enough, you ought to be able to correct for both paternal and maternal imprinting; the two types of imprinting have roughly similar complexity to each other. On the other hand, it might be harder to know the target state for sperm or for eggs, or harder to validate one or the other with trials.
- Using natural reproductive DNA ought to give diploid EC cells. If some GV method can avoid disrupting the epigenomic state, it likely would apply to both maternal and paternal DNA. But, this is far from guaranteed. For example, sperm DNA might be more difficult to work with than oocyte or embryonic DNA, because it is highly condensed.

Summary of genomic vectoring methods

A genomic vectoring method aims to create a cell that has a haploid or diploid complement of chromosomes that's been somehow steered (selected, modified). The following sections go into detail on several GV methods—the basic method, the obstacles to implementing it, and estimates of its genomic vectoring power. They are ordered roughly in terms of how hard it is to analyze the method's vectoring power. Simple embryo selection is the simplest to analyze coarsely; iterated recombinant selection is the most complicated selection GV method; and editing is not something I've worked on, so I'm deferring to others.

An annotated table of contents for the following Method sections:

- Simple embryo selection.
  - Make embryos, select the high-scoring one.
  - This is a weak GV method.
  - This is already being done.
- Gamete selection.
  - Select the high-scoring sperm and/or the high-scoring egg, and combine them.
  - This would be a fairly weak GV method.
  - This might become doable given non-destructive DNA sensing.
- Chromosome selection.
  - Select the high-scoring chromosomes from several cells and put them together in one cell.

  - Surprisingly strong. Chromosome selection on embryonic DNA from one couple would give over 6 SDs of raw selection power.
  - Requires some form of DNA sensing. The hardest part is getting the target chromosomes together in one cell, intact.
  - Has a good chance at bypassing the EC problem by using already-EC DNA, so is fairly *uncorrelated* with the success of almost all other reproductive GV methods.
- Iterated recombinant selection.
  - Induce diploid cells to divide meiotically into haploids, then combine the haploids to make new diploids, selecting high-scorers along the way.
  - This would be very strong GV.
  - This requires directly inducing meiosis, and requires an epigenomic correction method.
- Iterated multiplex CRISPR editing.
  - Apply multiple rounds of multiple CRISPR edits to lines of stem cells.
  - Given an epigenomic correction method, editing is very strong; without one, editing might be anywhere from weak to strong (by editing ESCs with epigenomic maintenance).

The following table summarizes the genomic vectoring power of the main GV methods. It's ordered roughly in terms of how feasible I think the method is—specifically, how soon I think it could be done, if someone seriously tried, soonest first. These estimates are uncertain and involve lots of modeling assumptions, but they're probably in the right ballpark.

Strength comparison of methods for strong human germline engineering

March 2025 | Berkeley Genomics Project | license: CC BY-SA 4.0 https://creativecommons.org/licenses/by-sa/4.0/

k	M(k)	M(2k)
1	0	0.56
2	0.56	1.03
4	1.03	1.42
10	1.53	1.87
23	1.93	2.22
300	2.87	3.09
20000	4.01	4.17

```
import numpy as np
def expected_max(test_runs=1_000, k=2):
  X = np.random.randn(test_runs, k)
  return np.mean(np.max(X, axis=1))"
```

M(k) means "the expected value, in SDs, of the largest of k samples from a standard gaussian"

10 embryo selection gives 1.53/√2=1 σ, 6 IQ points; haploid is .7 σ or 4 IQ

kman's editing estimates: https://www.lesswrong.com/posts/DfrSZaf3JC8vJdbZL/how-to-make-superbabies

for editing rows, column shows SDs for diploid editing, from kman's estimates based on genomics data that corresponds to a .35 PGS for IQ

methods organized in order of my very non-rigorous guess at how fast they *could* be doable if seriously tried

assuming R=.4 PGS

method name	prerequisites	formula for raw SDs	parameters	expected raw SDs	expected IQ points	notes
simple embryo selection	already doable	M(embryos)/√2	5 embryos	0.8 ± 0.6	5 ± 4	
			10 embryos	1.1 ± 0.5	6 ± 3	
			100 embryos	1.7 ± 0.4	10 ± 2	
			1000 embryos	2.2 ± 0.4	14 ± 2	
ESC CRISPR editing, small count	weak ESC imprinting maintenance (for a week, say); ESC SCNT works	6 + SDs × 15 from kman. diploid.	~15 edits?		14	0.56
sperm average-chromosome selection (+ simple embryo selection, 10 embryos)	sensing chromosome index chromosome transplantation sperm DNA must be mechanically robust enough	.7 + [M(donors)/√2] × (22/√46)	1 donor	0.7	4	
			2 donors	2	12	
			4 donors	3.1	19	
			10 donors	4.2	25	
			23 donors	5.1	31	
			300 donors	7.3	44	
			20,000 donors	9.9	59	
sperm recombinant chromosome selection (+ simple embryo selection, 10 embryos)	nondestructive DNA sensing, ~8 loci chromosome transplantation sperm DNA must be labellable sperm DNA must be mechanically robust enough	.7 + [1 + (M(donors)/√2)] × (22/√46)	1 donor	3.9	23	i.e. 1 father
			2 donors	5.2	31	
			4 donors	6.3	38	
			10 donors	7.5	45	
			23 donors	8.4	50	
			300 donors	10.5	63	
			20,000 donors	13.1	79	
....with replacement of k chromosomes		simulated	1 donor, 2 chr	4.6	28	
		donors here are *additional* to the main parent contributing the rest of the haploid	1 donor, 5 chr	5.2	31	
			4 donors, 2 chr	4.9	29	
			4 donors, 5 chr	5.7	34	
			23 donors, 1 chr	4.6	28	
			23 donors, 2 chr	5.1	31	
			23 donors, 4 chr	5.9	35	
			23 donors, 10 chr	7.7	46	
haploid iterated editing stem cells	spermatogonial stem cell induction/culturing, transplantation, artificial spermatogonia aren't outcompeted in vivo, GFP insert activated in sperm OR in vitro oogenesis OR in vitro spermatogenesis	4 + (SDs × 15 / 2)	15 edits		8	0.56
		from kman. 5-15 edits / round, perhaps. haploid (1/2 effect) + 4 points for embryo selection (monoclonal edited SSCs).	50 edits		13	1.26
			100 edits		18	1.81
			200 edits		23	2.49
			300 edits		26	2.96
			500 edits		31	3.63
			1000 edits		39	4.66
ESC chromosome selection	same requirements as gamete chr selection + weak maintenance + ESC SCNT; but DNA more accessible	same as sperm+egg recomb. chromosome selection; see below	1 couple	6.5	39	
sperm and egg recombinant chromosome selection	nondestructive DNA sensing, ~8 loci chromosome transplantation sperm + egg DNA must be labellable sperm + egg DNA must be mechanically robust enough	[1 + (M(donors)/√2)] × (44/√46)	1 donor each	6.5	39	yes really! i think.
			2 donors each	9.1	55	
			4 donors each	11.2	67	
			10 donors each	13.5	81	
		"1 donor each" means one progenitor of each haploid, i.e. a single couple	23 donors each	15.3	92	
			300 donors each	19.7	118	do not do this
			20,000 donors each	24.9	149	do not do this
....with replacement of k chromosomes		simulated	1 donor, 2 chr	7.9	47	
		donors here are *additional* to the main parent contributing the rest of the haploid	1 donor, 5 chr	9	54	
			4 donors, 2 chr	8.4	50	
			4 donors, 5 chr	10	60	
			23 donors, 1 chr	7.8	47	first chromosome replacement gives +8 IQ points
		replacement and chr selection for each haploid separately	23 donors, 2 chr	8.8	53	
			23 donors, 4 chr	10.4	62	
			23 donors, 10 chr	14	84	
in vivo spermatogonial editing	delivery, low off-target rates, good efficiency, fluorescent tagging that shows in sperm; edited spermatogonia have to stay alive	6 + (SDs × 15 / 2) from kman. haploid (1/2 effect) + 6 IQ for full embryo selection	~50 edits?		15	1.26
iterated editing ESCs	maintain naive ESC epigenetic state, including imprinting, for months; ESC SCNT if imprinting not perfectly maintained, epigenetic editing OR full diploid epigenomic correction	6 + SDs × 15	15 edits		14	0.56
		from kman. 5-15 edits / round, perhaps. diploid (full effect)	50 edits		25	1.26
			100 edits		33	1.81
			200 edits		43	2.49
			300 edits		50	2.96
			500 edits		60	3.63
			1000 edits		76	4.66
haploid chromosome selection	nondestructive DNA sensing, 1 locus chromosome transplantation spermatogonial stem cell induction, transplantation, artificial spermatogonia aren't outcompeted in vivo, GFP insert activated in sperm OR in vitro oogenesis OR in vitro spermatogenesis	.7 + M(2 × donors) × (22/√46)	1 donor	2.5	15	
			2 donors	4	24	
			4 donors	5.3	32	
			10 donors	6.8	41	
		(haploid)	23 donors	7.9	47	
			300 donors	10.7	64	
			20,000 donors	14.2	85	
iterated meiotic selection	direct induction of meiosis full diploid epigenomic correction (OR one of IVO / IVS / iSSC transplantation, for half the effect)	M(2 × couples) × √rounds × (44/√46)	1 round, 1 couple	3.6	22	
			1 round, 2 couples	6.7	40	
			1 round, 4 couples	9.2	55	
			1 round, 10 couples	12.1	73	
		hard to determine bounds. estimates here assume rigid segments, but allow a round of chromosome selection at the end. without that assumption these numbers are optimistic by a lot, but probably still give a good general sense.	2 rounds, 1 couple	5.1	31	
			2 rounds, 2 couples	9.4	56	
			2 rounds, 4 couples	13	78	
			2 rounds, 10 couples	17.2	103	do not do this
			3 rounds, 1 couple	6.3	38	
			3 rounds, 2 couples	11.6	70	
			3 rounds, 4 couples	16	96	
			3 rounds, 10 couples	21	126	do not do this
			4 rounds, 1 couple	7.3	44	
			4 rounds, 2 couples	13.4	80	
			4 rounds, 4 couples	18.4	110	do not do this
			4 rounds, 10 couples	24.3	146	do not do this
			5 rounds, 1 couple	8.1	49	
			6 rounds, 1 couple	8.9	53	
			7 rounds, 1 couple	9.6	58	
			8 rounds, 1 couple	10.3	62	

(Link to spreadsheet, which you could inspect, copy, and modify: https://docs.google.com/spreadsheets/d/1AaBs7d772q_UgReIMCYM3m1-Jc3QHvcBhfg47gsK0ug/edit?usp=sharing. I have found (and epigenomically corrected) several errors in drafts of this spreadsheet, so please be on the lookout.)

The order statistics $M(k)$ given at the top, which are used in many of the power estimates for selection methods, can be checked with this snippet:

```
import numpy as np
def expected_max(test_runs=10_000_000, k=2):
    X = np.random.randn(test_runs, k)
    return np.mean(np.max(X, axis=1))
```

All the IQ point numbers are additional over the mean population score q, and assume that all input genomes are drawn from the population at random, with a standard normal distribution centered at q. The mean q can be whatever you want, though at the very high end there is a ceiling on IQ (it's not known where or with what distribution).

This table is only a summary. More analysis is given in the following sections. There are many combinations of GV methods and EC-making methods that aren't specifically covered, though in general combining GV methods would significantly increase genomic vectoring power.

Except for simple embryo selection, all of these methods are speculative, at least in humans. No one has done them yet, and furthermore they all require scientific and technological understanding that hasn't been achieved yet. They will likely have many unforeseen obstacles (though will also in some cases be easier than expected).

All the discussion below focuses on making a single child. See "Selection makes similar children" for a potential issue with this assumption.

Method: Simple embryo selection

Method:

- Do in vitro fertilization (IVF) to get some embryos.
- Sequence the embryos by taking a few cells from each.
- Pick the highest-scoring embryo to try to implant.

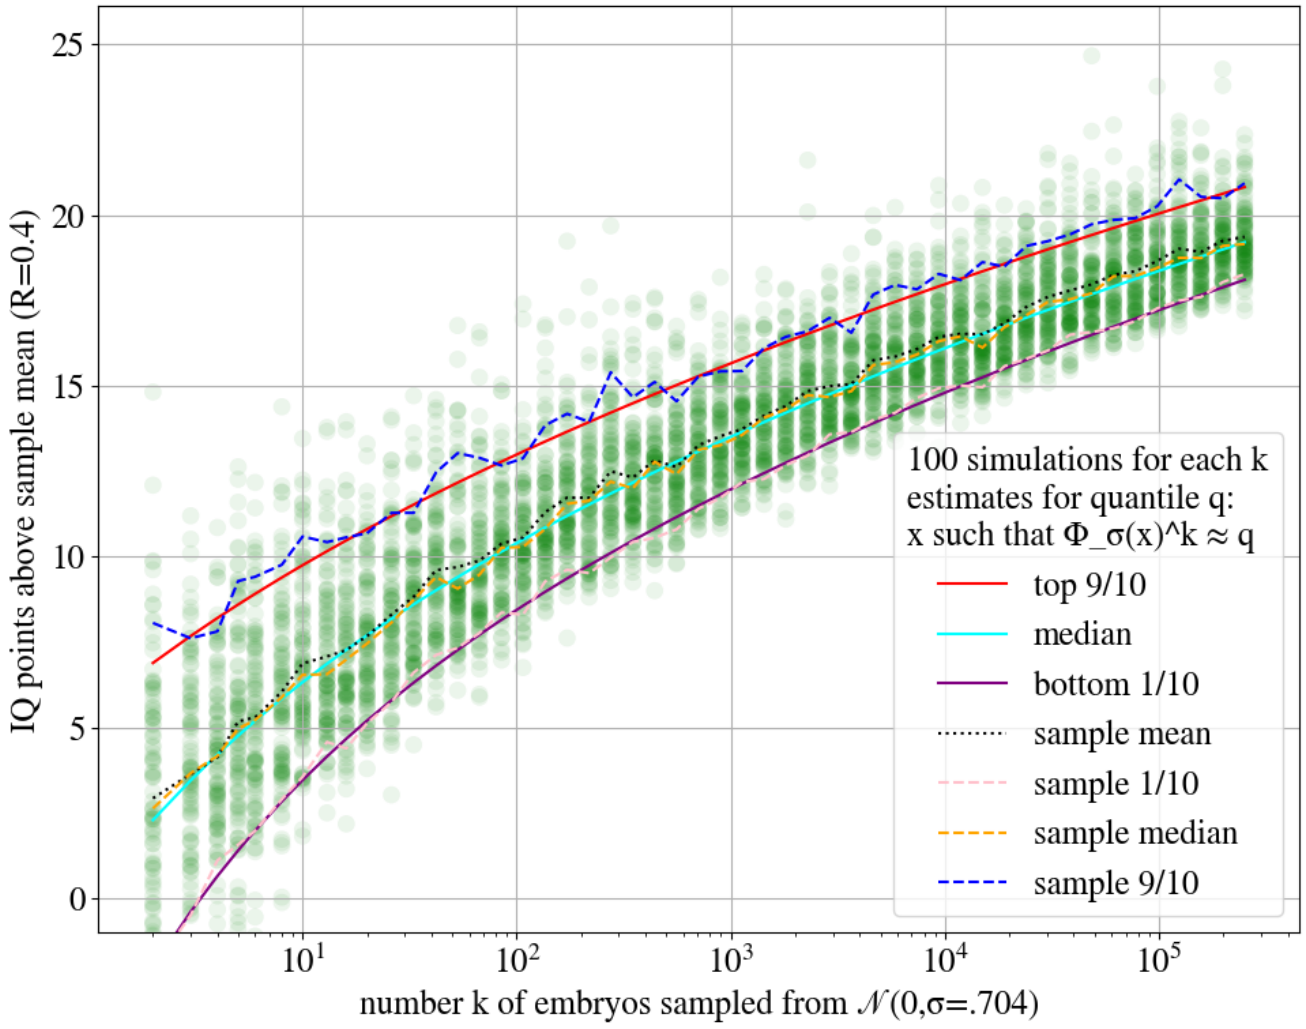
Simple embryo selection is already being done; see for example Orchid, Lifeview, Heliospect, and GenEmbryonics.

The power of simple embryo selection

The effect of embryo selection is severely limited by two factors:

- IVF produces only a small number of achievable births, i.e. eggs produced during superovulation which are successfully fertilized, develop normally enough in vitro, and then successfully implant and gestate. An older women might obtain only a few or even no embryos; with multiple superovulation rounds and good reproductive health, a young woman might optimistically hope to have one or two hundred achievable births to select from.
- The variance in the genomes of embryos is half the population variance (the variance of a random person's genome).

Together, these effects severely limit the practical effects. Assuming a PGS that correlates at 0.4 with IQ, we get:



This graph uses a simple model, where the embryos for one couple are drawn from a gaussian with variance roughly .5 (more exactly, .496). To give some numbers:

all estimates from the median of $\mathbb{E}_\sigma[\mathcal{N}(0,\sigma)]$k, which gives the raw selection power
the ± ranges describe the 80% range of outcomes between the bottom 1/10 and top 1/10 outcome.
(the number given is half that range.)

number (k) of embryos sampled. eX means 10^X

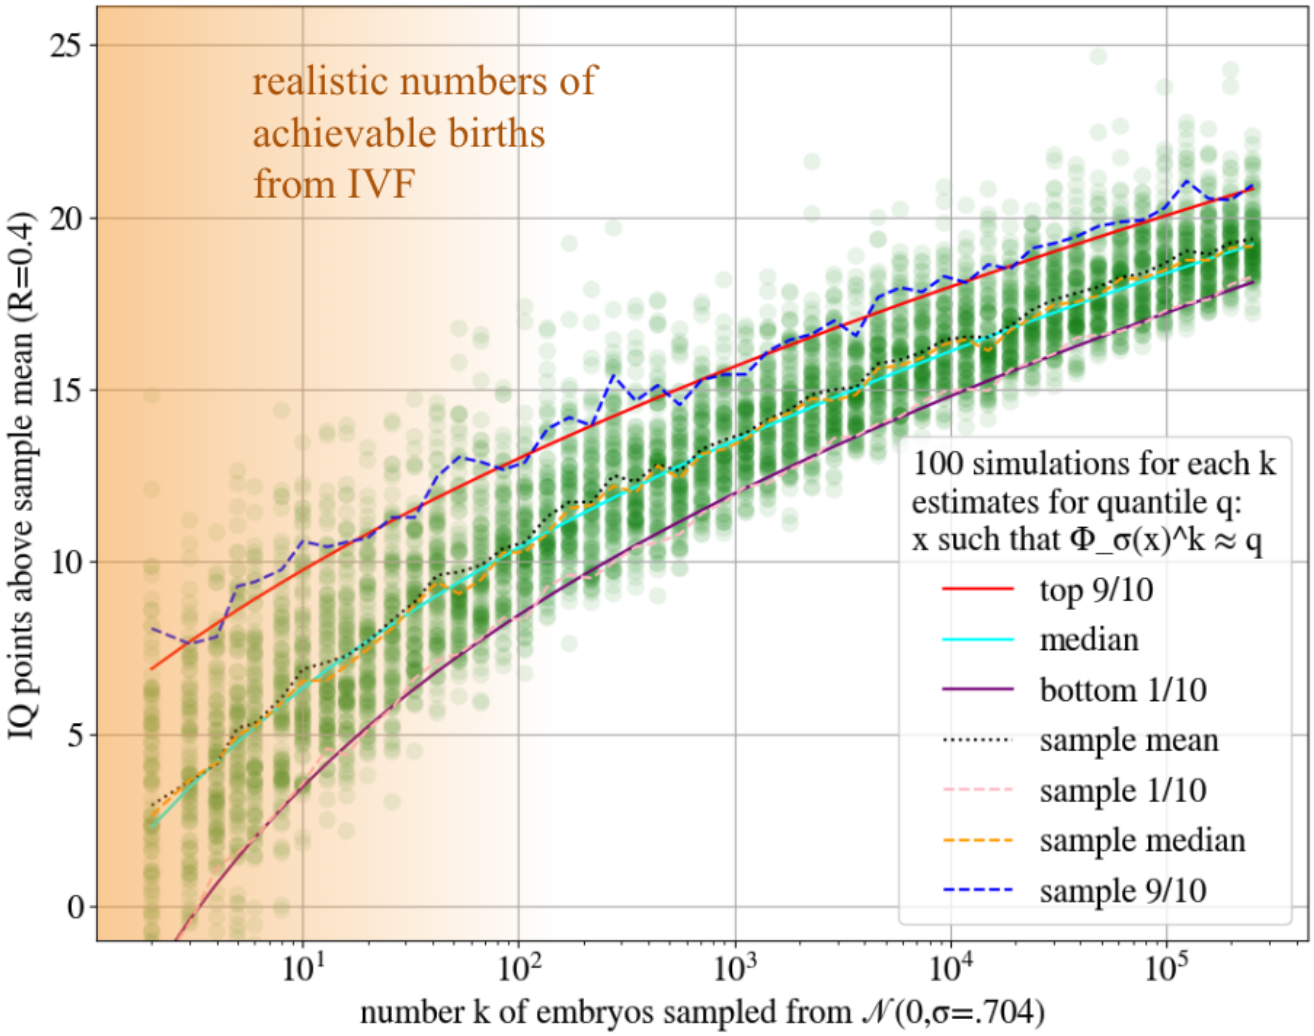
	$\mathbb{E}_\sigma[\mathcal{N}(0,\sigma)]$k	effective IQ SDs	IQ points
1	0.00 ±0.90	0.00 ±0.36	0 ±5
2	0.38 ±0.74	0.15 ±0.30	2 ±4
3	0.57 ±0.67	0.23 ±0.27	3 ±4
4	0.70 ±0.63	0.28 ±0.25	4 ±4
5	0.79 ±0.60	0.32 ±0.24	5 ±4
6	0.86 ±0.58	0.35 ±0.23	5 ±3
7	0.92 ±0.56	0.37 ±0.23	6 ±3
8	0.97 ±0.55	0.39 ±0.22	6 ±3
9	1.01 ±0.54	0.41 ±0.22	6 ±3
10	1.05 ±0.53	0.42 ±0.21	6 ±3
e1.5	1.42 ±0.45	0.57 ±0.18	9 ±3
e2.0	1.73 ±0.40	0.69 ±0.16	10 ±2
e2.5	2.01 ±0.38	0.80 ±0.15	12 ±2
e3.0	2.25 ±0.36	0.90 ±0.14	14 ±2
e3.5	2.48 ±0.35	0.99 ±0.14	15 ±2
e4.0	2.69 ±0.35	1.08 ±0.14	16 ±2
e4.5	2.89 ±0.34	1.15 ±0.14	17 ±2
e5.0	3.07 ±0.35	1.23 ±0.14	18 ±2
e5.5	3.24 ±0.35	1.30 ±0.14	19 ±2
e6.0	3.41 ±0.35	1.36 ±0.14	20 ±2
e6.5	3.57 ±0.36	1.43 ±0.14	21 ±2
e7.0	3.72 ±0.37	1.49 ±0.15	22 ±2

cells show SDs (or IQ points)
IQ points assumes a PGS with R of .4
effective IQ SDs = (SDs of raw selection power) × .4 (effective SDs / raw SD)
IQ points = (effective IQ SDs) × 15 (IQ points / effective SD)

This table shows the gains on IQ if you select the top-scoring embryo out of k embryos, where k is shown in the leftmost column. (In what follows, in the context of IVF, k really means achievable births.) The leftmost column is k; the next column shows raw SDs of selection power. The next column shows effective-SDs or trait-SDs for IQ. Raw selection power is just "how uncommon is this embryo"; the effective selection power or trait selection power is "how uncommon is this embryo, on the scale of this trait". If the PGS for IQ correlates with IQ at 0.4, then a 1 SD exceptional embryo on our PGS will be around 0.4 SDs exceptional on actual IQ.

(Different couples have different expected gains, though not due to their average IQ. The distribution used in the table above integrates over the uncertainty over the couple's embryo-variance σ, and the values given are the median with an error range spanning 10th percentile to 90th percentile. Also, the median values can be fairly easily approximately calculated as follows: take the max-of-k distribution for a standard gaussian, given in "Samples to standard deviations"; and then divide the given number by $\sqrt{2}$, since the variance of one couple's embryos is 1/2 the variance of the general population. E.g. a rarity of 1 in 10^6 corresponds to 4.8 SDs; $4.8/\sqrt{2} \approx 3.4$; which is what the above table gives for e6.0. But the table also gives a sense for the range of likely outcomes, which is harder to estimate. See the appendix "Detailed estimation of embryo selection".)

We quickly hit a big wall when sampling more and more. Going from 1 to 100 embryos gives +10 IQ points; going from 100 to 1,000,000 also gives +10 IQ points. We don't get very far with realistic achievable births with current methods:



Super-duper-ovulation

Several research groups are attempting to work out a method for in vitro oogenesis—making eggs artificially by taking stem cells and stimulating in them the cellular processes that turn them into eggs. If one of those groups succeeds, you could generate large numbers of eggs, fertilize them, and then apply simple embryo selection.

It might also be possible to harvest many eggs via ovarian tissue extraction, stimulation of dormant follicles[131], and in vitro maturation. Ovarian tissue cryopreservation is a treatment practiced to restore fertility after chemotherapy, so we know that ovarian tissue can produce eggs after having been extracted[132]. Human oocytes can be fully matured from the adult stage (arrested in prophase I; antral follicle) to an oocyte, in vitro[133]. However, the current practical efficiency of these methods is not great[134].

In principle this method might be able to produce many natural eggs in order to make many more embryos than is possible via superovulation. However, normally as follicles grow and oocytes mature in preparation to be ovulated, they suppress each other until one is dominant and is ovulated. This ovarian follicle dominance might stop you from getting, say, hundreds or thousands of eggs from extracted ovarian tissue. In the most

[131]Kawamura, Kazuhiro, Yuan Cheng, Nao Suzuki, Masashi Deguchi, Yorino Sato, Seido Takae, Chi-hong Ho, et al. 'Hippo Signaling Disruption and Akt Stimulation of Ovarian Follicles for Infertility Treatment'. Proceedings of the National Academy of Sciences 110, no. 43 (22 October 2013): 17474–79. https://doi.org/10.1073/pnas.1312830110.

[132]Poirot, Catherine, Anne Fortin, Nathalie Dhédin, Pauline Brice, Gérard Socié, Jean-Marc Lacorte, Jean-Paul Akakpo, et al. 'Post-Transplant Outcome of Ovarian Tissue Cryopreserved after Chemotherapy in Hematologic Malignancies'. Haematologica 104, no. 8 (August 2019): e360–63. https://doi.org/10.3324/haematol.2018.211094.

[133]Hatırnaz, Şafak, Barış Ata, Ebru Saynur Hatırnaz, Michael Haim Dahan, Samer Tannus, Justin Tan, and Seang Lin Tan. 'Oocyte in Vitro Maturation: A Sytematic Review'. Turkish Journal of Obstetrics and Gynecology 15, no. 2 (June 2018): 112–25. https://doi.org/10.4274/tjod.23911.

[134]Nikiforov, Dmitry, Cheng Junping, Jesus Cadenas, Vallari Shukla, Robert Blanshard, Susanne Elisabeth Pors, Stine Gry Kristensen, et al. 'Improving the Maturation Rate of Human Oocytes Collected Ex Vivo during the Cryopreservation of Ovarian Tissue'. Journal of Assisted Reproduction and Genetics 37, no. 4 (1 April 2020): 891–904. https://doi.org/10.1007/s10815-020-01724-7.

successful case, this would be limited to under half a million eggs, as adult women have at most that many oocytes with primordial follicles.

One could imagine getting tens of thousands of eggs, or more. This would be a 50% to 100% improvement over simple embryo selection, which is still not very strong.

You would want to use some extremely cheap screening method, to avoid paying millions of dollars in sequencing costs. Batched small SNP arrays might be cheap enough, or if not, some more complicated method might work; the appendix "Cheap DNA segment sensing" discusses this question. If you used FISH, you'd just pick the embryo whose biopsy was the brightest. (Or rather, pick the brightest 10 or 50, and then whole genome sequence them and pick the highest scoring one.)

Method: Gamete selection

Single gamete selection

Method:

1. Collect a bunch of gametes (either a bunch of sperm or a bunch of eggs).
2. Somehow find a small subset of those gametes that are higher scoring.
3. Combine those gametes with complementary gametes, producing embryos.
4. Sequence the embryos and pick the highest-scoring one to try to implant.

Gamete selection is mostly hypothetical

The basic problem is that all high-information DNA analysis (sequencing, SNP genotyping) is destructive—to sequence a cell, you destroy it, amplify the DNA, and then apply a lossy detector of some sort. Further, by default there's no way to culture a gamete. Both sperm and eggs are in a kind of stasis; their DNA is largely epigenomically silenced, so they don't produce the proteins that would be needed to mitose.

In theory one could do "entanglement sequencing", wherein you capture the four meiotic "grandchildren" of a single gametogonium (progenitor stem cell that produces gametes). This is impractical as far as I know, because both oogenesis and spermatogenesis are complex and all but require a bunch of surrounding tissue. Spermatogenesis in particular involves a complicated process where many partially-differentiated progenitor cells stay tightly connected as they pass through the seminiferous tubules. Oocytes might be more feasible to indirectly sequence, e.g. by looking at the polar body emitted after false fertilization.

In any case, there's little point in doing oocyte selection unless you either also have sperm selection, or you're more constrained on processing embryos than on producing oocytes. Otherwise you're just going to fertilize all your oocytes with a random sperm anyway and then sequence it, so it doesn't really matter if you know the oocyte's haploid genome beforehand.

But sperm selection might be doable

What you'd need is some method for non-destructively determining enough about a sperm's haploid genome, that you can sort many thousands or millions of them. I don't know of a method that clearly ought to work.

Something like FISH might work. You would label several variants that appear in dad's paternal chromosome but not in his maternal chromosome, or vice versa, depending on which segment is higher scoring. See the appendix "Cheap DNA segment sensing". Then you'd pick the brightest or most multicolored sperm, e.g. using FACS.

The power of single gamete selection

How many SDs of selection pressure do you get from single gamete selection?

Variance is additive, i.e. $\mathrm{Var}(X+Y) = \mathrm{Var}(X) + \mathrm{Var}(X)$. The variance of an embryo is the sum of the variances of the gametes from each person. So the variance among the gametes sampled from one person is half the variance of the embryos of a couple. In terms of the effects of selection, everything scales in terms of SDs σ. The SD of a gaussian with variance σ^2 is σ, i.e. the square root of the variance. So the gametes drawn from one person, which have a variance $1/2$ that of embryos drawn from a couple containing that person, will have a standard deviation of $\sigma_g = \sigma_e/\sqrt{2}$.

We can use this to convert from the numbers for simple embryo selection to numbers for gamete selection. If we *only* did gamete selection, taking the top-scoring of k gametes, and didn't select the resulting embryos after fertilization, the selection powers would be $1/\sqrt{2}$ times whatever the power would have been for selecting k embryos.

Double gamete selection

Suppose that sperm selection works, as described above. We'd then combine sperm selection with embryo selection. Once you have one or several top-scoring sperm, you'd fertilize some eggs to obtain embryos; sequence them; and pick the highest-scoring ones.

We can think of this as essentially being double gamete selection, which in general is the following method:

1. Get a bunch of sperm, pick the highest-scoring one.
2. Get a bunch of eggs, pick the highest-scoring one.
3. Fertilize the egg with the sperm.

Doing sperm selection followed by embryo selection is not quite the same power as double gamete selection. Instead of pairing the very best sperm with the very best egg, you pair the top, say, 20 best sperm with 20 random eggs, and then pick the best; you probably didn't pair the best sperm with the best egg.

On the other hand, if you select the top 20 from millions of sperm, they'll mostly be around what you expected to get from the top 1. Then when you select your embryo, you add on the power of selecting the best egg. So the loss compared to actual double gamete selection is not necessarily very noticeable. (Technically, you do better than best egg plus top 20 sperm, because maybe the second best egg was paired with the top sperm, and together they are higher-scoring than the best egg paired with another sperm.)

Thus we can roughly view sperm selection followed by embryo selection as being double gamete selection. Further, suppose we have in vitro oogenesis. Then it would become worth it to do oocyte selection as well as sperm selection, getting the full power of double gamete selection.

The power of double gamete selection

If we abstract away those details and consider the mathematical situation, how many SDs of (raw) selection power does double gamete selection give?

Using the discussion above, we can give a naive answer. (See "The power of selection" for a more convoluted discussion[135].) If we combine the single best of k sperm with the single best of k eggs, we get something that is $2/\sqrt{2} = \sqrt{2} \approx 1.41$ times more SDs than the single best of k embryos from the same couple. Call the latter value $M_e(k)$, meaning "expected maximum of k samples from the e(mbryo) distribution", and likewise $M_g(k)$ for g(ametes). So we have

$$M_g(k) + M_g(k) = 2M_g(k) = 2M_e(k)/\sqrt{2} = \sqrt{2}M_e(k)$$

and in general,

$$M_g(k_1) + M_g(k_1) = \frac{M_e(k_1) + M_e(k_1)}{\sqrt{2}}$$

So for example, suppose you can use in vitro oogenesis to make 1,000 eggs, and you can select from 10,000 sperm. Then you look up the values M_e in the table above, reproduced here:

[135] Benson-Tilsen, Tsvi. 'The Power of Selection', 9 August 2022. https://tsvibt.blogspot.com/2022/08/the-power-of-selection.html.

all estimates from the median of $\mathbb{E}_\sigma[\mathcal{N}(0,\sigma)]$k, which gives the raw selection power
the ± ranges describe the 80% range of outcomes between the bottom 1/10 and top 1/10 outcome.
(the number given is half that range.)

number (k) of embryos sampled. eX means 10^X

	$\mathbb{E}_\sigma[\mathcal{N}(0,\sigma)]$k	effective IQ SDs	IQ points
1	0.00 ±0.90	0.00 ±0.36	0 ±5
2	0.38 ±0.74	0.15 ±0.30	2 ±4
3	0.57 ±0.67	0.23 ±0.27	3 ±4
4	0.70 ±0.63	0.28 ±0.25	4 ±4
5	0.79 ±0.60	0.32 ±0.24	5 ±4
6	0.86 ±0.58	0.35 ±0.23	5 ±3
7	0.92 ±0.56	0.37 ±0.23	6 ±3
8	0.97 ±0.55	0.39 ±0.22	6 ±3
9	1.01 ±0.54	0.41 ±0.22	6 ±3
10	1.05 ±0.53	0.42 ±0.21	6 ±3
e1.5	1.42 ±0.45	0.57 ±0.18	9 ±3
e2.0	1.73 ±0.40	0.69 ±0.16	10 ±2
e2.5	2.01 ±0.38	0.80 ±0.15	12 ±2
e3.0	2.25 ±0.36	0.90 ±0.14	14 ±2
e3.5	2.48 ±0.35	0.99 ±0.14	15 ±2
e4.0	2.69 ±0.35	1.08 ±0.14	16 ±2
e4.5	2.89 ±0.34	1.15 ±0.14	17 ±2
e5.0	3.07 ±0.35	1.23 ±0.14	18 ±2
e5.5	3.24 ±0.35	1.30 ±0.14	19 ±2
e6.0	3.41 ±0.35	1.36 ±0.14	20 ±2
e6.5	3.57 ±0.36	1.43 ±0.14	21 ±2
e7.0	3.72 ±0.37	1.49 ±0.15	22 ±2

cells show SDs (or IQ points)
IQ points assumes a PGS with R of .4
effective IQ SDs = (SDs of raw selection power) × .4 (effective SDs / raw SD)
IQ points = (effective IQ SDs) × 15 (IQ points / effective SD)

We have $M_e(10^3) \approx 2.2$ and $M_e(10^4) \approx 2.7$. So we get

$$M_g(10^3) + M_g(10^4) \approx \frac{2.2 + 2.7}{\sqrt{2}} \approx 3.5$$

SDs of raw selection power. Multiplying by the PGS strength of .4, we get 1.4 SDs of IQ, or 21 IQ points. Pretty cool, though not quite strong GV.

The limits of extreme gamete selection

This simple model will break down. The distribution of gamete scores is not really a gaussian. There's some distribution induced over single chromosomes by the crossover process; then 23 of these variables are added together to get the haploid genome of a gamete. Adding together a bunch of variables kinda makes a gaussian, but not exactly.

The upshot is that outside of a few SDs, gamete selection basically hits a wall, defined by whatever the best crossovers for every chromosome are. I think this wall is somewhere around 9 SDs of raw power? I'm not sure. That guess is based on the best crossover being about +1.3 standard deviations (see the appendix "Best crossover"); applying that to each of 46 chromosomes, each with 1/46 of the variance, gives $1.5 \times 46/\sqrt{46} \approx 9$. But there are other relevant variables not accounted for here. For example, there may be quite a large variation in how many crossovers there are in a given gamete. E.g. a 99th percentile crossover-count sperm might have as many as 37 crossovers, meaning many chromosomes will have a few crossovers[136]. That raises the ceiling a lot. In particular I'm not sure when the gaussianity assumption

[136]Bell, Avery Davis, Curtis J. Mello, James Nemesh, Sara A. Brumbaugh, Alec Wysoker, and Steven A. McCarroll. 'Insights into Variation in Meiosis from 31,228 Human Sperm Genomes'. Nature 583, no. 7815 (July 2020): 259–64. https://doi.org/10.

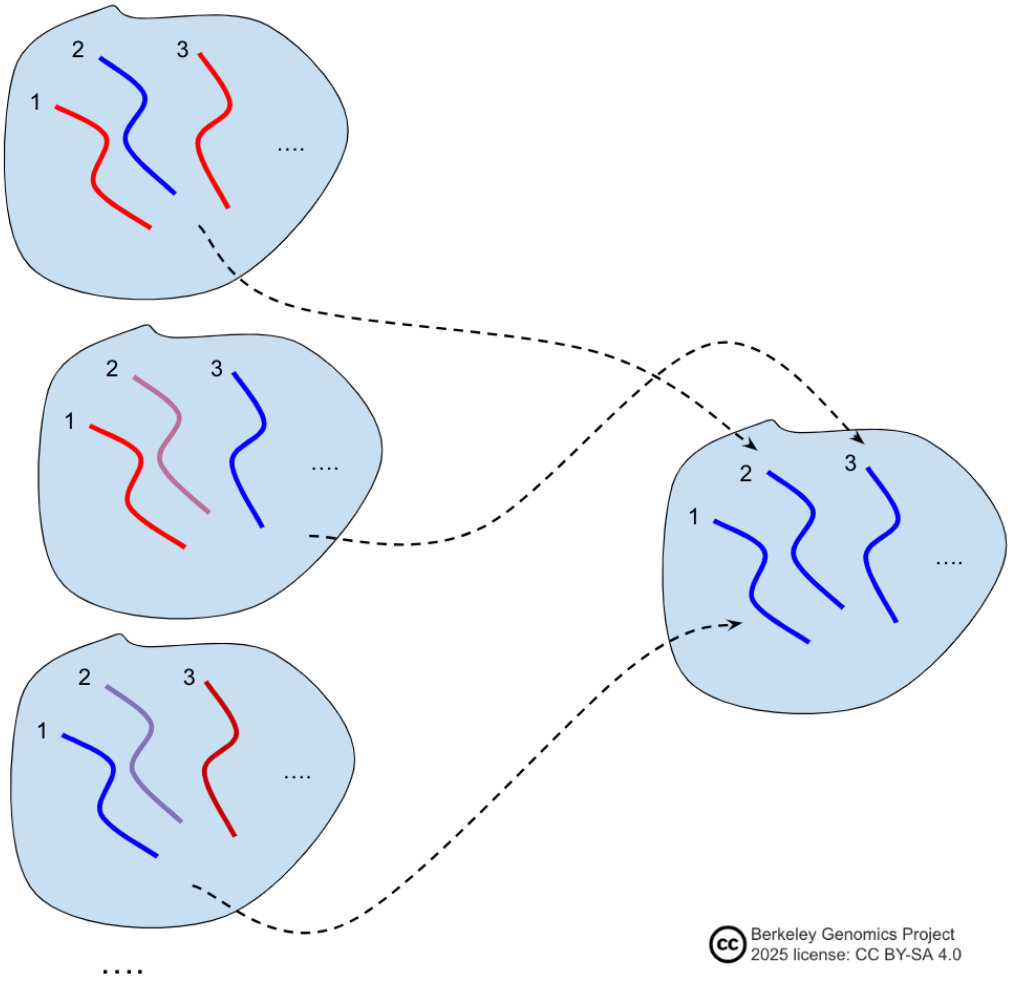
materially breaks down. However, at a wild guess I'd think you could get to +6 SDs of raw selection power, or +2.5 SDs of IQ, i.e. around 50 IQ points. Not bad, if you can select from huge amounts of gametes.

Method: Chromosome selection

Method:

1. Take one or more cells.
2. Get enough information about their genomes to identify target chromosomes.
3. Somehow arrange that target chromosomes from different cells make into one cell.
4. Use that cell to make a baby.

An illustration:



The earliest discussion of chromosome selection I'm aware of is this comment, referenced by Branwen[137].

Two key points about chromosome selection:

- Both average chromosome selection with many donors, as well as diploid recombinant chromosome selection, are truly strong GV methods. See below.
- In theory, some chromosome selection methods don't disrupt epigenomic states.
  - Some hypothetical chromosome selection methods just move DNA around physically for a short period of time, rather than e.g. culturing cells for several weeks or months. In theory, if applied to natural reproductive DNA (oocytes, sperm, embryonic stem cells), this should preserve the DNA's epigenomic state.
  - Thus, in theory, chromosome selection can bypass the need for epigenomic correction, which is a major obstacle for most reproductive GV methods. This doesn't mean that chromosome selection is easy; rather, it means that the success of developing chromosome selection is fairly *decorrelated*

1038/s41586-020-2347-0.

[137] Gwern Branwen. 'Embryo Selection For Intelligence', 22 January 2016. https://gwern.net/embryo-selection.

with the success of developing other methods, because chromosome selection has a different set of obstacles.
- E.g.: you could apply chromosome selection to sperm DNA.
- E.g.: you could apply chromosome selection to SSCs, and then transplant the frankenSSCs back into a man's testes.
- E.g.: in theory, you could apply chromosome selection to fully grown germinal vesicle primary oocytes. You'd put the chromosomes into a donor germinal vesicle primary oocyte in antral follicle, and then transplant the frankenOocyte into adult ovaries or apply in vitro maturation. (The point of this would be to avoid wasting scarce fully mature oocytes. You'd have to start diakinesis somehow, because the diplotene chromosomes are still connected.)

Implementation of chromosome selection

There are several kinds of method for chromosome selection. See a list of possible methods here: "Chromosome selection". In general, chromosome selection requires some way to identify chromosomes. See cheap DNA segment sensing. Some of the methods in the list linked above don't require direct DNA sensing; e.g. targeted chromosome destruction followed by electrofusion might work.

Two example methods:

Mechanical manipulation

One chromosome selection method that is potentially especially feasible is using FISH to identify chromosomes, combined with mechanical transplantation to a cell. You label the chromosomes, lyse the cells, somehow grab the target chromosome, and then inject it into your target cell.

A major potential obstacle to this kind of method is that it could mechanically break the DNA. There's research about using flow cytometry to study chromosomes, including methods to keep chromosomes intact[138]. So maybe it is feasible.

A great benefit of this kind of method is that it ought not to interfere much with the epigenomic state of the chromosome, assuming that the medium is not disruptive.

An obstacle here is identifying the chromosome. FISH may work, but as discussed in "An obstacle with sperm", sperm DNA may be inaccessible. On the other hand, ESC DNA ought to be accessible for FISH (but could be more vulnerable to mechanical breaking). Therefore this method could be used in an EC bypassing way! This should require minimal culturing of the cell—just enough to get many tries, and then to finally verify the integrity of the final zygote—so loss of imprinting should be small, hopefully.

Another major obstacle is grabbing the target chromosome. One possibility is to use a micromanipulator. However, chromosomes are very small and difficult to see in a normal microscope. It might be doable but difficult with fluorescent labeling. Another possibility is to use FISH and FACS, although FACS might fragment the chromosomes due to shear forces. Another possibility is using magnetic beads attached to the DNA, and then magnetically grabbing them. However, magnetic beads are large, so they might hurt the DNA or might collect dozens of homologous chromosomes.

For sperm DNA, there's a further issue: The DNA is highly compacted and maybe tangled together. That means there's a good chance that it's not feasible to straightforwardly grab a chromosome. It might be feasible to decondense sperm chromatin with chemical treatments[139] and then grab some chromosomes.

Immature oocytes are arrested in diplotene stage, before diakinesis. That means homologous chromosomes are still connected to their homologs by chiasmata (recombination points), and sister (copy) chromatids are still connected to each other at their centers (forming little Xs). Possibly you could stimulate diakinesis and then perform a kind of average recombinant selection on crossed-over sister-chromatid pairs to form a secondary-oocyte-like cell. But this is probably too complicated.

Mature oocytes, arrested in metaphase II, may be more amenable to chromosome extraction. You can artificially activate an MII oocyte as though it were fertilized, and it will proceed with meiosis. After a

[138] Doležel, Jaroslav, Sergio Lucretti, István Molnár, Petr Cápal, and Debora Giorgi. 'Chromosome Analysis and Sorting'. Cytometry Part A 99, no. 4 (2021): 328–42. https://doi.org/10.1002/cyto.a.24324.

[139] Chankitisakul, Vibuntita, Nutthee Am-In, Theerawat Tharasanit, Tamas Somfai, Takashi Nagai, and Mongkol Techakumphu. 'Sperm Pretreatment with Dithiothreitol Increases Male Pronucleus Formation Rates After Intracytoplasmic Sperm Injection (ICSI) in Swamp Buffalo Oocytes'. Journal of Reproduction and Development 59, no. 1 (2013): 66–71. https://doi.org/10.1262/jrd.2012-104.

short while the chromosomes should be nicely separated, while still condensed as chromatin.

These problems shouldn't apply to ESCs.

Whole cell fusion

See "Whole cell fusion". (A.k.a. "poor man's chromosome selection".)

The basic idea is to just fuse two diploid cells together, e.g. via electrofusion, to form a tetraploid cell. Then, somehow, you induce the tetraploid cells to spontaneously correct their ploidy, dividing into diploid cells. One can thus iteratedly shuffle chromosomes around, applying selection pressure to keep the target chromosomes, until you get a diploid cell with most of the target chromosomes.

In other words, iterated whole cell fusion and tetraploid correction is a kind of iterated selection, like iterated meiotic selection but without the meiosis.

As a GV method, this might be doable right now, though it might be pretty inconvenient. Mouse tetraploid cells injected into an embryo were found to sometimes spontaneously correct to diploidy, sometimes giving euploid cells[140]. Maybe one could induce this in vitro.

I assume this method would not maintain a naive ESC epigenomic state, so it would separately require epigenomic correction after producing a GV cell. If tetraploidy reduction were feasible, the whole cell fusion protocol for chromosome selection could be tested soon in non-human species.

The power of simple chromosome selection

In simple chromosome selection, the input cell lines are independent of each other—they're from different people, and don't share abnormally much DNA. Examples:

- You start with several people, and take somatic cells from each.
- You start with several embryos, each one from a different couple.
- You start with one person, and take just one sample.

In any case, the chromosomes can be roughly modeled as just being random samples from a gaussian with variance 1/46 the general population variance, because variance is additive.

Selecting a haploid genome

In the simplest case, suppose that we construct a cell that will play the role of a sperm, by assembling a haploid subset of one parent's own whole genome (that is, without meiotic recombination). We assume that this DNA will fertilize a randomly sampled oocyte. How strong is this protocol?

If we focus on chromosomes at index 1, what happens? We select the best of 2 samples from a gaussian. Very roughly, this gives a result 0.5 SDs above the mean (more precisely, 0.56), in terms of just this chromosome's SDs. We do this 22 times, once for each autosome, and a chromosome-SD is about $1/\sqrt{46}$ a normal SD. So we get roughly

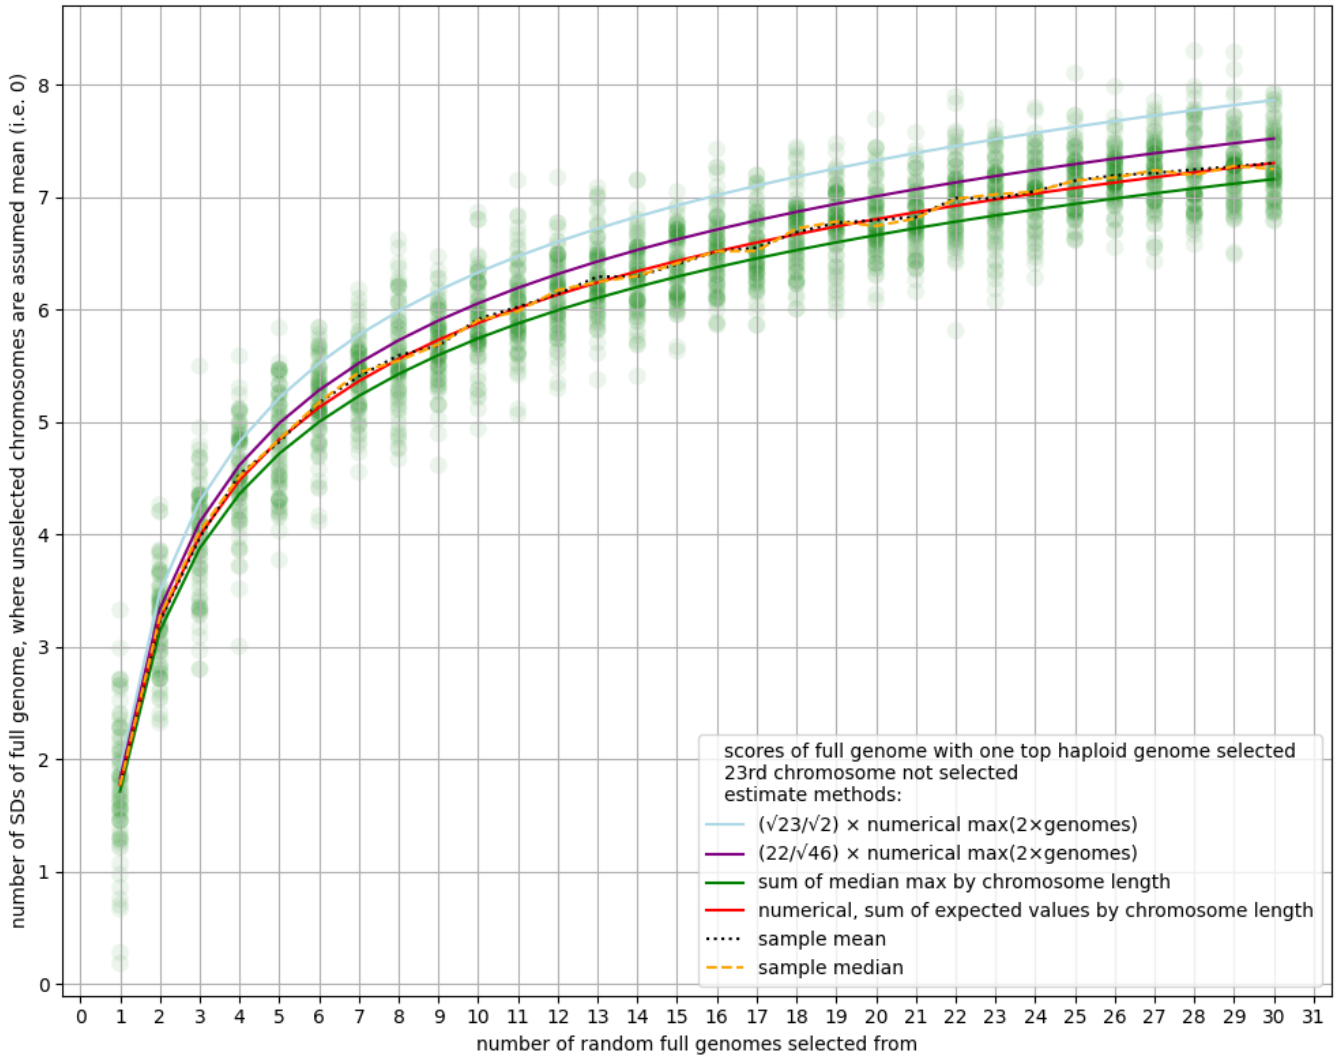
$$.5 \times 22/\sqrt{46} \approx 1.6$$

SDs. This is fairly weak; if we applied all this raw selection power to IQ, we'd get $.4 \times 1.6 \approx .6$ SDs, or 9 IQ points. Adding in a round of embryo selection gets perhaps another .7 raw SDS, for a total of 13 IQ points.

("Hold on,", you might be saying, "chromosomes have quite different lengths; the longest is 5 times as long as the shortest! Doesn't this greatly decrease the selection power, since variance is pretty concentrated compared to if the chromosomes were all of equal length?" That's a good question, but no, it makes very little difference. See the appendix "Variation in chromosome length".)

What would happen if you took a larger set of partial parents, and selected the best chromosomes taken from any of them? If we construct a haploid this way, we get:

[140] Frade, João, Shoma Nakagawa, Paola Cortes, Umberto di Vicino, Neus Romo, Frederic Lluis, and Maria Pia Cosma. 'Controlled Ploidy Reduction of Pluripotent 4n Cells Generates 2n Cells during Mouse Embryo Development'. Science Advances 5, no. 10 (16 October 2019): eaax4199. https://doi.org/10.1126/sciadv.aax4199.



(This and following graphs are from "The power of selection", which gives more discussion and code for the graphs[141].)

At the bottom left, we see 1.7 SDs for a single genome (best 1 of 2 at each index), roughly as our estimate above said.

Looking further to the right, at 27 on the x-axis, suppose a woman chooses 27 people. Then we sequence those 27 people, and make a haploid from their collective top-scoring chromosome at each index. Then she has a baby gotten from that haploid combined with her egg. Reading the y-axis (the red curve), we see about 7.3 SDs. Combined with a round of embryo selection, that's about 8 SDs of raw selection power. If it's all applied to IQ, that's around 50 IQ points. (If instead she took from 10,000 people, she could get about 80 IQ points for her child, above some combination of the mean of those people's and her own genomic IQ.)

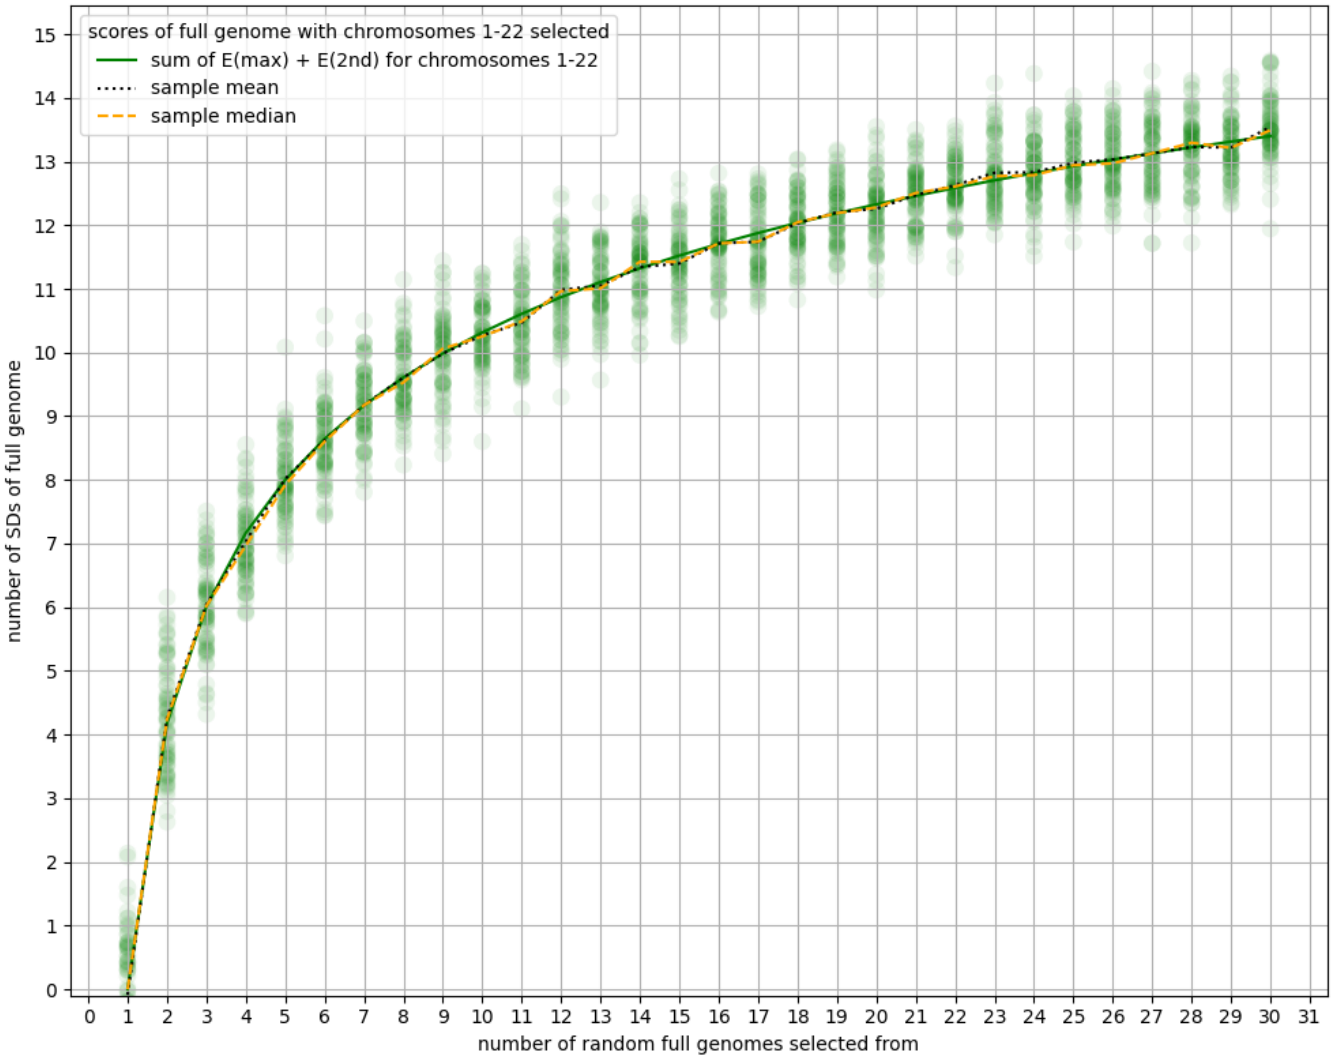
Selecting a diploid genome

Suppose instead that we select the entire genome. There are several possible permutations of this, and I won't go through them all in detail.

E.g., a natural question is, what if you separately create two haploids, one from each parent or one from each set of parents, and combine them? Then you basically add the gains from each. If it's two parents, you get about 3.4 raw SDs, or about 20 IQ points. If it's many parents, you could do quite well. E.g. with 4 people for each of the two haploids, you can get 9 raw SDs, or about 55 IQ points.

Another kind of question is, what if you create the whole diploid together, by selecting across all the chromosomes? Now you're selecting the top 2 out of some set of chromosomes. Here are the results from that:

[141] Benson-Tilsen, Tsvi. 'The Power of Selection', 9 August 2022. https://tsvibt.blogspot.com/2022/08/the-power-of-selection.html.



The results are quite strong. (Though, apparently, not much stronger than doing haploids separately—I think this is understandable as "top 1 of k" being not that far from "top 1 of $2k$".) Four people can get 7 SDs of raw selection power, or 40 IQ points. (Suppose the same four people want to have several children this way, but without just using the same chromosomes, and keeping the total amount of each person's DNA propagation roughly the same. I don't know how strong this would be, but at a wild guess, we can imagine that the four people could have several children, each with +30 IQ, who are genetically related to all four as genetic grandchildren.)

Another multi-parent arrangement:

- One could replace a small portion of a haploid with chromosomes from other people. For example, we could replace the lowest-scoring quarter of a haploid's chromosomes.
- As a very rough estimate, doing this with, say, around 50 chromosome donors might increase the total raw selection power by an average of $3.5/\sqrt{46} \approx .5$ SDs per chromosome. If both parents did this with 6 chromosomes, the result would be an additional 6 SDs of raw selection power, or 35 IQ points.
- They'd be genetically related to their children at a closeness halfway between a parent and a grandparent. (One imagines their friends doing similarly, so that no one loses much in total DNA propagation, if that's something they mind. This protocol might, however, potentially have the effect of making their children look noticeably different from themselves—though presumably less different than grandchildren vs. grandparents.)

Selecting an epigenomically correct haploid

Some applications of chromosome selection could bypass the EC problem by selecting from DNA that's already in an epigenomically correct state, rather than applying an EC-making method.

For example, we could start with some embryos from several couples, and then assemble an embryo by selecting a haploid genome of chromosomes with *maternal* imprinting, and another haploid genome of chromosomes with *paternal* imprinting. When we do this, we're constrained in which chromosomes to pick for each haploid. Each embryo only contains one chromosome 1 with paternal imprinting, and so on.

In this example, the selection power for k starting embryos can be computed by just looking up the power for $k/2$ starting diploids. That's because the set of available chromosomes to be chosen from is halved.

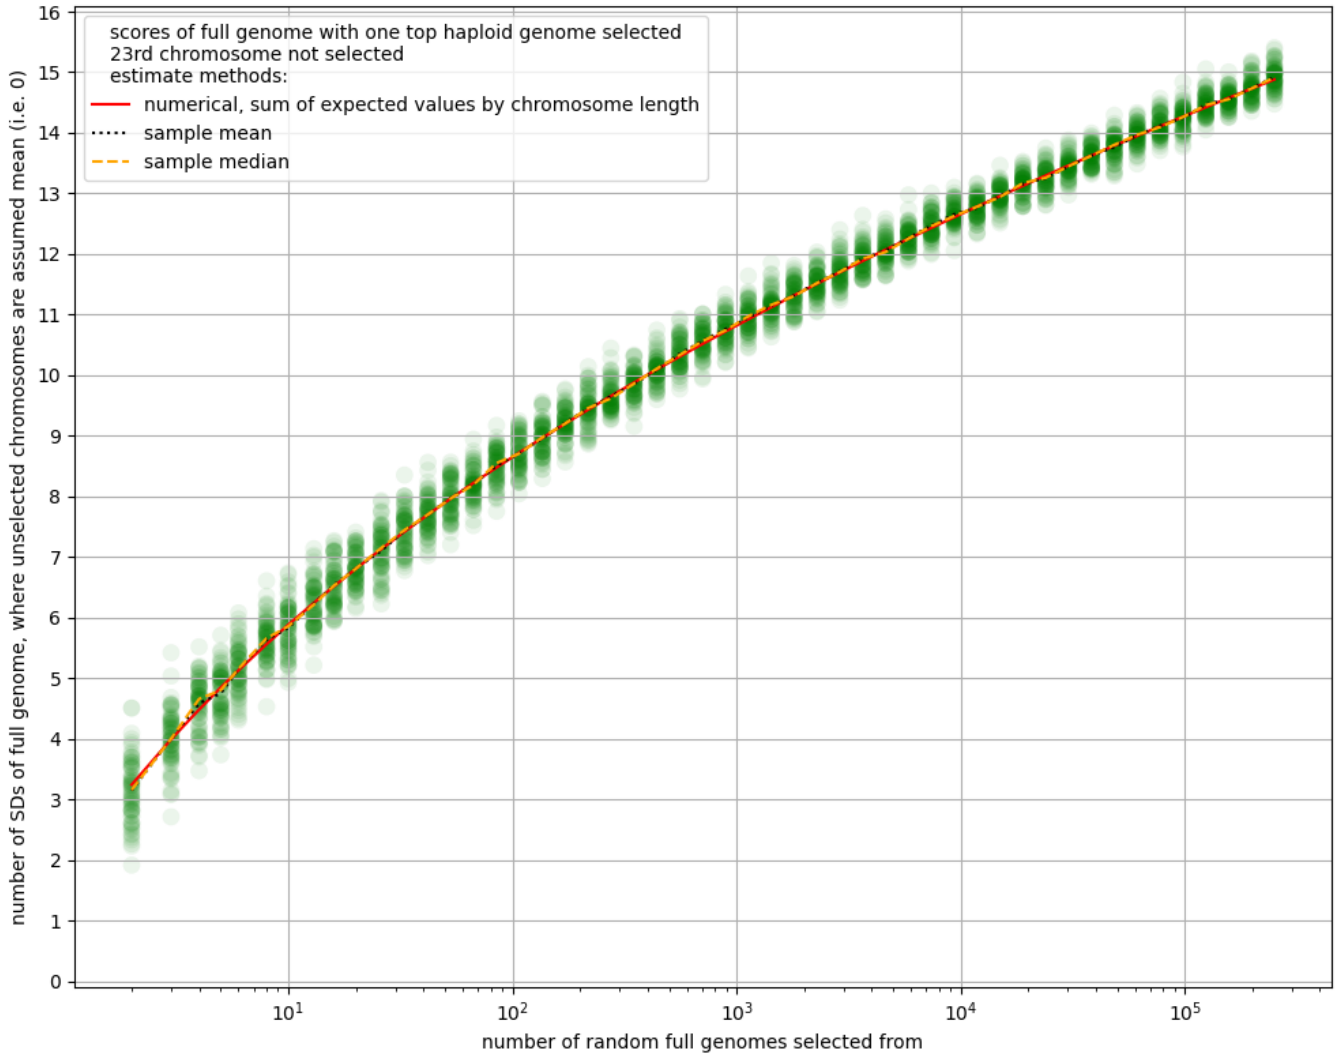
(This example is rather contrived, and is mainly just for purposes of understanding the math. In real life you'd more likely be doing a form of recombinant chromosome selection; see the next subsection.)

Selecting average chromosomes

It might be relatively easy to distinguish chromosomes by index, without being able to distinguish anything else about them. See "Sensing chromosome index".

In that case, given a started set of many donors, you could perform chromosome selection at each index by selecting the person with the highest average score of their two chromosomes at that index. The power of this form of selection would essentially be like chromosome selection, except all the power would be divided by $\sqrt{2}$. (The variance of the average of two chromosomes is 1/2 the variance of a single chromosome.) You also have to have twice the number of input genomes to be comparable, because for each index you only get one average-chromosome per genome, whereas you get two chromosomes per genome.

As one example, you could do average-chromosome selection on sperm DNA from many men to create a super sperm. How powerful is this? We can compute this using the analysis in "Selecting a haploid genome". Here's a version of the graph in that section, but with many more input genomes:



So suppose we have 20,000 men. That's the equivalent of 10,000 whole genomes to select from (because we only get one sample of an average-chromosome per genome). The graph shows about 12.5 SDs of power from chromosome selection. But we have to adjust because the variance of average-chromosomes is less. We get: $12.5/\sqrt{2} \approx 8.8$ SDs. That's actually pretty strong!

If instead we take a few hundred men, we get about $9/\sqrt{2} \approx 6.4$ SDs. Combined with simple embryo selection, that's about 7 raw SDs, or about 40 IQ points. Not bad either.

Can this really be right? We can double check by computing the same thing a bit more directly. With

around 200 samples from a gaussian, the max is around 2.6 SDs. We have $2.6 \times 22/\sqrt{2 \times 46} \approx 6$. With this estimate we still get around 40 IQ points.

As discussed above, similar gains could be gotten while still having a strongly genetically related father. For example, you could use more donors, and replace only the bottom-average-scoring half or quarter of the father's chromosomes.

This has a notable implication: **If it's possible to nondestructively isolate chromosomes by index, then it's possible to make very intelligent children.** For example, if using a centrifuge to separate many chromosomes by mass works to make bands of different chromosome index, and there's some way to isolate single ones (e.g. into liposomes), one could do this.

Selecting a diploid source of haploids

Suppose you extract SSCs from one or more men, apply chromosome selection to make a frankenSSC clone, and transplant it back into a man's testes. How much selection power do you get?

If we ignore problems with homozygous SSCs, you can do at least as well as simple haploid chromosome selection: you just double up on the best chromosome at each index. If there are several donors, you can use the top 2, which should fix homozygosity issues.

But there's a further question: can you do better?

You can't increase the expectation of the resulting haploids, of course. But crossovers between two copies of the same chromosome literally do nothing, so the all-top diploid produces haploids with no variance.

You can get greater variance by using some non-top chromosomes. Is this strictly a loss? No: crossover chromosomes can be higher scoring than the highest scoring single starting chromosome; see below.

So the plan would be to pick pairs of chromosomes (from multiple donors of course) that increase variance in the resulting haploids, while still being high average scoring. Then you'll do gamete selection somehow, or simply do embryo selection after fertilization with haploids from the frankenSSCs.

How do you pick which chromosomes to put in the SSC, to maximize the expected maximum embryo score from this procedure? What is the resulting expected maximum embryo score? I don't know. If you figure it out or run simulations let me know.

The power of recombinant chromosome selection

In recombinant chromosome selection, the input cell lines are not independent of each other. Rather, they are meiotic child cells of some set of donor cells. Examples:

- You start with one or more couples, and take several embryos from each.
- You start with several gametes from one person.
- You start with several people, and take several gametes from each.
- You start with one or more people, and use in vitro meiosis to produce several haploid cells from each person.

The advantage of recombinant over simple chromosome selection is that there is more variation in the scores of chromosomes that come from meiotic DNA recombination.

In terms of implementation, these methods are basically the same as simple chromosome selection. The main difference is that, in order to take advantage of the variance due to recombination, we have to sense multiple points on each chromosome. See the appendix "Cheap DNA segment sensing".

I haven't done the work to get a precise estimate of the power of recombinant chromosome selection in various forms. But see the following subsections. There are several complicating aspects:

- The distribution that meiotic recombination induces on the scores of the child chromosomes is quite importantly not a gaussian.
- The power depends to some extent how many (2, 3, or many) points on each chromosome we can distinguish; or in other words, how closely we can locate the crossover points.
- The power depends on how many crossover points occur. While the average for sperm is around 1.1, it varies by quite a bit. Depending on the numbers, you might be able to get a bump in power.
- Crossover points don't happen at random points.
- See the appendix "Best crossover".

One-donor recombinant chromosome selection

To simplify further, we can analyze the case where we assume (unrealistically):

- We select chromosomes from the gametes / haploids from just one person.
- Every chromosome has exactly 1 crossover point.
- The crossover point is uniformly distributed throughout the chromosome.
- Each chromosome is a sequence of many small gaussian samples, whose variances add up to the variance of the whole chromosome.

See the appendix "Best crossover" for some discussion of what the highest-scoring meiotic child chromosome will be. The upshot is that the very best chromosome with 1 crossover will be about 1.25 SDs above mean (in terms of SDs of that chromosome), and the very best with 2 crossovers will be about 1.75 SDs.

To make a more reasonable assumption than optimality, we can consider that we're able to label each chromosome at about 6 places. This gives about 1 SD of selection power for each chromosome; in other words, about twice as powerful as simple chromosome selection on two chromosomes. If we do this for a whole haploid genome, we get $1 \times 22/\sqrt{46} \approx 3.2$ raw SDs, or about 24 IQ points after a round of embryo selection.

My guess is that this is actually an *underestimate*, because there are often two or three crossovers in a chromosome. A gold standard feasibility check would be to get single-cell genotyping data from a sperm sample and then compute the results of chromosome selection on that set of chromosomes using real PGSes.

Many-donor recombinant chromosome selection

Now consider multiple donors. If we did average-chromosome selection, we'd get some result. If instead we do recombinant chromosome selection, what do we get?

I think we can get a lower bound on selection power as follows. First we sample each donor's average chromosome score. Then we pick the highest average score. Then we get a good recombination of the two chromosomes. The recombination will be about 1 chromosome-SD higher-scoring than the average score. In other words, we maximize over the average-chromosomes of a given index; then we add 1. We might do better by comparing the best recombinations between donors, rather than just comparing the average scores (though obtaining the average score is far cheaper, unless we reliably get near the best recombination from one person's haploid sample).

In the case of one donor, average-chromosome selection gives +0, and then we add 1, giving the same as the above estimate.

In the case of 2 donors, we have $1 + .56/\sqrt{2} \approx 1.4$ as the corrected power at each chromosome. Then we pass that through each of 22 chromosomes, and add .7 with an embryo selection stage at the end. This gives 5.2 SDs of raw selection power. Someone spot-check the methodology here! The formula I'm using:

$$\text{embryo selection} + \frac{\left(\text{boost avg. to best recombination} + \frac{\text{best 1 of } k}{SD \text{ scale of average}}\right) \times 22 \text{ autosomes}}{SD \text{ scale of chromosomes}}$$

For 2 donors, and embryo selection on 10 embryos (but correcting down by $\times 1/\sqrt{2}$ because only selecting the other haploid), this gives

$$.7 + \frac{\left(1 + \frac{0.56}{\sqrt{2}}\right) \times 22}{\sqrt{46}} = 5.2$$

Which is rather a lot. That's +31 IQ points with one mother, two fathers, and sperm chromosome selection.

Double recombinant chromosome selection

If oocyte DNA is tough enough to handle the mechanical stress, and sperm DNA can be labeled well enough, and we can isolate chromosomes, then we could do recombinant chromosome selection on both gametes.

In a spherical cow model, that doubles the power. The situation is complicated by the fact that oocytes might be in quite limited supply. However, the space of crossovers is quite limited; it's a two-dimensional space (for the positions of two crossovers, as oocyte chromosomes often have), or one-dimensional if we're

only optimizing over the 1-crossover chromosomes. With, say, 10 oocytes, we should get in the ballpark of the best chromosome. Since the above estimate is already conservative, a rough estimate of double haploid recombinant chromosome selection is simply double the power (but without much benefit from a round of embryo selection afterward): around 6.5 raw SDs, or around 40 IQ points.

I'll repeat that:

> **With one mother, one father, and recombinant chromosome selection, you get more than +40 IQ points.**

Oocyte chromosome selection has the unfortunate effect of destroying many of the available oocytes. Since IVF embryos quite often don't make it all the way to birth (due to trouble fertilizing, morphological abnormalities, or normal miscarriage). This is likely too much of a cost for many mothers, but it would be technically feasible, assuming that the chromosome extraction method can work reliably and with single cells (which may be very difficult). IVO would remove this obstacle by making oocytes with different crossovers abundant.

Another way to avoid this obstacle, and give more than one chance to extract the chromosomes, is ESC chromosome selection. You grow the zygote into an early-stage embryo, say for just two weeks. Then you label the DNA somehow, and take a few cells from embryos with target chromosomes. From those cells you extract the target chromosomes. This shouldn't destroy many embryos. However, there's an additional challenge: identifying which chromosome is paternal and which is maternal. This may not be hard, e.g. with FISH targeted to alleles that only one of the parents has.

Fractional haploid donation

Another permutation of chromosome selection is this: instead of making a haploid by selecting all the highest-scoring chromosomes from any of several donors, you have a main haploid progenitor (presumably a member of a couple looking to have a baby), and then several external donors. The main haploid will have some especially low-scoring chromosomes, which will be replaced by the best chromosome from the external donors's haploid chromosomes (assuming that maximum is higher than that the main haploid's version). This would allow the main haploid progenitor to be more closely genetically related to their child than a grandparent would be, though still less genetically related than a parent.

This table shows the selection power of different numbers of donors and chromosomes replaced:

first row shows number of external donors used
first column shows number of chromosomes replaced

chromosomes replaced

	0	1	2	4	10	23
1	6.7	7.4	7.5	7.6	7.8	7.8
2	6.7	7.9	8.2	8.4	8.6	8.8
3	6.8	8.3	8.7	9.0	9.3	9.6
4	6.7	8.7	9.1	9.5	10.0	10.4
5	6.7	9.0	9.5	10.0	10.6	11.1
6	6.8	9.2	9.8	10.4	11.2	11.8
7	6.8	9.3	10.1	10.8	11.7	12.4
10	6.8	9.7	10.8	11.8	13.0	14.0
12	6.8	9.7	11.1	12.4	13.7	14.9

cells show SDs for the whole diploid genome
each haploid separately has [column] external donors
up to [row] many chromosomes are replaced by the best from donors
replacements are in order of most improvement
no replacement if no improvement (rare except for 1 donor, 10 or 12 chrs)
best crossovers found, allowing dividing chromosomes into 5 segments
each cell simulated with 500 samples

In row k, column d, this table shows the result of doing recombinant chromosome selection on both haploids separately, with k replacements each. For one parent's haploid contribution, we sample d donor haploid genomes. For each chromosome index, we consider the maximum recombined chromosome from the d donors at that index, and look at its improvement over the parent's chromosome at that index. For the k largest improvements, we replace the parent chromosome with the donor chromosome.

To get the values for this protocol applied to only one haploid, divide by 2 (and add .7 for a round of embryo selection).

Look at the rightmost column. With 23 donors, replacing just 3 chromosomes in each parent's haploid is nearly an additional 3 SDs above double recombinant chromosome selection (shown in the left column, under 0). At 9.6 raw SDs, this protocol would add 58 IQ points. Even further:

first row shows number of external donors used
first column shows number of chromosomes replaced

chromosomes replaced

	0	1	10	300
1	6.7	7.4	7.7	8.2
2	6.8	7.9	8.6	9.4
3	6.8	8.3	9.3	10.4
4	6.8	8.7	9.9	11.3
5	6.7	8.9	10.5	12.4
6	6.9	9.3	11.2	13.3
7	6.8	9.3	11.7	14.2
10	6.8	9.5	13.0	16.5
12	6.7	9.8	13.7	18.0

cells show SDs for the whole diploid genome
each haploid separately has [column] external donors
up to [row] many chromosomes are replaced by the best from donors
replacements are in order of most improvement
no replacement if no improvement (rare except for 1 donor, 10 or 12 chrs)
best crossovers found, allowing dividing chromosomes into 5 segments
each cell simulated with 50 samples

(There might be more noise in this table due to the lower sample size.)

Here's the code for simulating random genomes (modeled as subdivided gaussians): https://gist.github.com/tsvibt/89cec1b4fd7d54be04ba34a0059bcc4c

Here's the code for generating the specific data in this table: https://gist.github.com/tsvibt/3a3104d5e3530306c02ee7cbd2cf7ad7

Method: Iterated recombinant selection

Method:

1. Take a population of cells.
2. Create haploid recombinant child cells via meiosis.
3. Combine haploids together to make diploids.
4. Repeat 1—3; dovetail sequencing some of the cells and selectively amplifying some of them.
5. At the end, apply some EC-making method to make a baby.

Because iterated recombinant selection involves meiosis, an epigenomically correct state can't be maintained. Besides the reprogramming done during meiosis itself, maternal and paternal segments of homologous chromosomes would end up in a single crossed-over chromosome; the result is epigenomically correct neither as a paternal nor as a maternal chromosome. Therefore iterated recombinant selection has to be followed by an EC-making method such as in vitro gametogenesis.

That said, iterated recombinant selection is extremely powerful. Because it can progressively interleave smaller and smaller segments of the starting chromosomes, the achievable selection power is, in practice, unbounded. In other words, you could get so many raw SDs of selection power that it would literally be unsafe to apply them all to known traits. See the bullet point "Traits outside the regime of adaptedness" in the perils list.

In practice, the question about iterated recombinant selection is about the costs. My guess, based on preliminary analysis and simulations, is that you can get quite large effects—over 10 raw SDs—for a few hundred thousand dollars near the beginning, or less. (The costs should then follow an innovation curve.) See the appendix "The costs of iterated meiotic selection" for more detail.

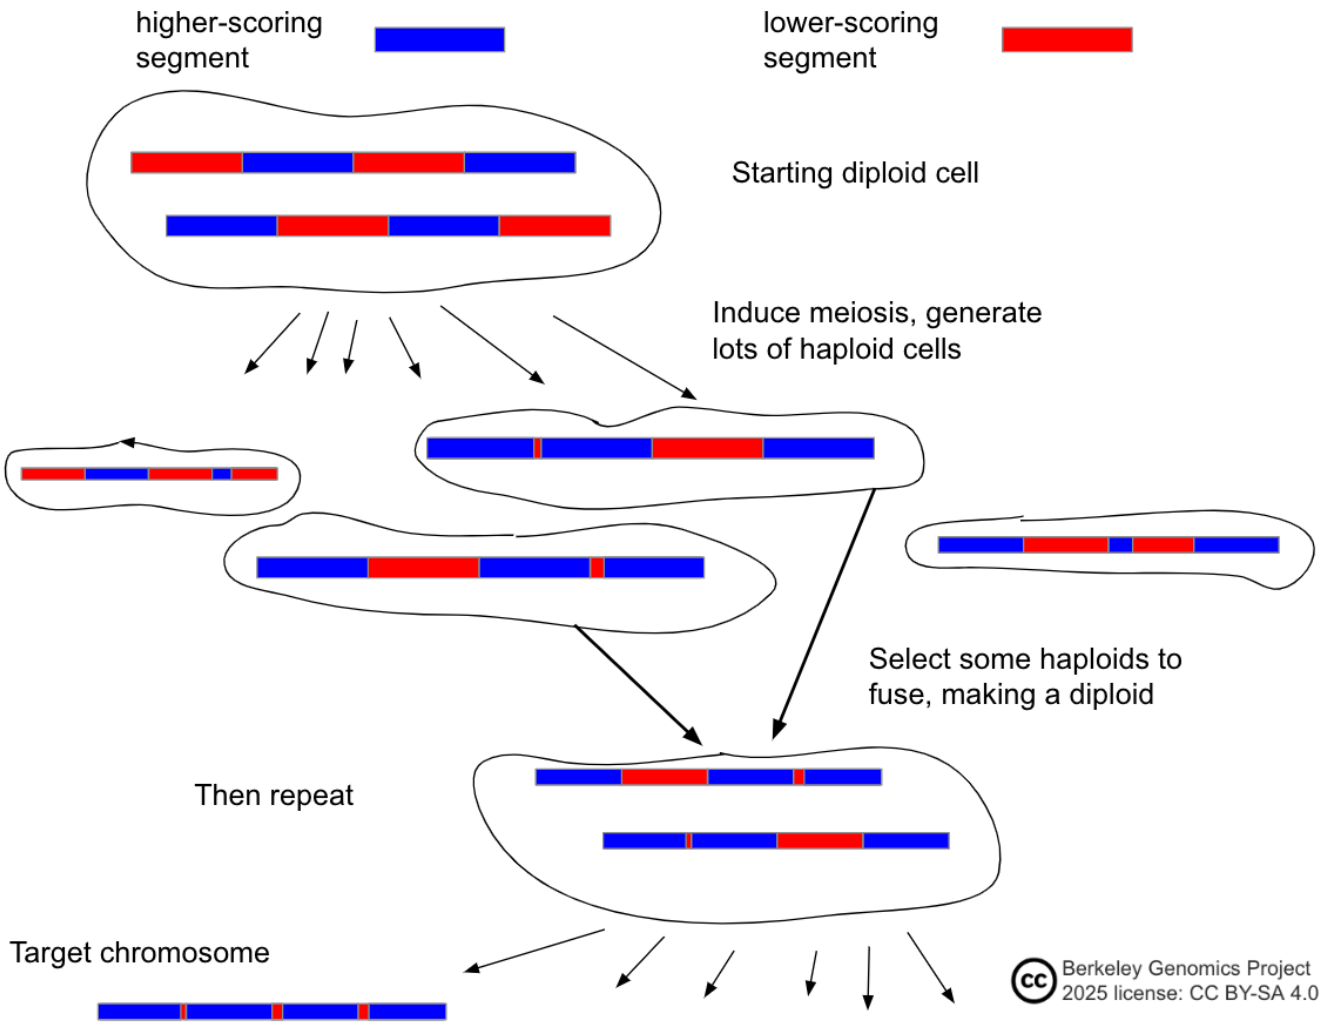
Iterated embryo selection

See "History of Iterated Embryo Selection", Branwen 2019[142].

In iterated embryo selection, the haploid cells are created via in vitro spermatogenesis and in vitro oogenesis. The resulting gametes are used to create embryos, which are grown in vitro; then the cycle repeats.

Iterated meiotic selection

The first description of IMS I'm aware of is "Meiosis is all you need", Metacelcus, 2022[143]. IMS is roughly the same as iterated embryo selection, except that instead of trying to recapitulate the whole reproductive cycle including spermatogenesis, oogenesis, fertilization, and embryonic growth, you directly induce meiosis in stem cells. Instead of fertilization, you simply fuse two cells through electrofusion[144]. An illustration:



As Metacelsus points out, there are three major advantages to IMS over iterated embryo selection:

- It's more powerful to select haploids rather than diploids.
  - I'm not sure exactly how much, because the situation is complex, but a very rough rule of thumb is that haploid selection might go about $\sqrt{2} \approx 1.4$ times faster than diploid selection. (This is based on the simple case of selecting from a gaussian X vs. selecting from each of two i.i.d. Y each with half the variance of X.)

[142] Gwern Branwen. 'History of Iterated Embryo Selection', 18 January 2019. https://gwern.net/ies-history.

[143] Metacelsus. 'Meiosis Is All You Need'. Substack newsletter. De Novo (blog), 30 June 2022. https://denovo.substack.com/p/meiosis-is-all-you-need.

[144] Trontelj, Katja, Marko Ušaj, and Damijan Miklavčič. 'Cell Electrofusion Visualized with Fluorescence Microscopy'. Journal of Visualized Experiments: JoVE, no. 41 (1 July 2010): 1991. https://doi.org/10.3791/1991.

- Meiosis is much faster than gametogenesis. In humans:
  - Spermatogenesis: very roughly 3 months.
  - Oogenesis: at least 6 months naturally; could maybe do it faster in vitro.
  - Meiosis: 3 weeks, maybe possible to do it in 2. That's roughly an order of magnitude faster.
  - Someone who is good at clocks please help me budget this.
- IMS doesn't require IVO and IVS.
  - Gametogenesis involves both meiosis, and also epigenomic reprogramming and other morphological changes. Recapitulating the whole process may be significantly more complex than merely causing a cell to undergo meiosis. IMS only calls for inducing meiosis, which may be simpler to discover methods to do.
  - However, IMS does still require some form of EC-making at the end (which could e.g. be IVO).

A fourth advantage over iterated embryo selection: IMS does not involve the creation of embryos. Haploid cells that haven't been epigenomically corrected are not functional as gametes, and fusing them to form a diploid stem cell does not create a cell that could naturally grow into a embryo. This should alleviate some life-protecting concerns about killing many embryos in vitro, which would happen in iterated embryo selection.

The main challenge to implementing IMS is to induce cells to undergo meiosis in vitro. This has not been achieved scalably and with appreciable efficiency in human cells. See the discussion of meiosis in the earlier section "IVG as an epigenomic correction method".

In the human body, meiosis only happens in the germline—spermatocytes and oocytes. So it might be that the easiest way to get meiosis is to invent in vitro gametogenesis. On the other hand, the bar is lower for meiosis compared to gametogenesis; no EC-making is needed, for purposes of IMS. Natural gametogenesis takes a long time because there are several other processes happening (epigenomic reprogramming, oocyte growth), and because gametogenesis is tied to the whole organism's development. So it's plausible that inducing meiosis is easier than full IVG, and could be applied for much faster meiosis. As a suggestive observation: during the growth of oocytes in primordial follicles, part of the function of cumulus granulosa cells is to *prevent* the oocyte from resuming meiosis[145]. Maybe *removing* cumulus cells when imprinting has been established, even if the oocyte isn't fully grown, would speed up in vitro meiosis compared to natural oogenesis.

One obstacle to full IMS is the culturability of the haploids. To do selection over haploids, you'd have to proliferate them monoclonally so you can take a sample and destructively genotype them, and then selectively amplify some clones. Culturability may or may not come for free. Some strange sort of human haploid embryonic stem cells have been cultured[146]. (Even without culturability, the cycle speedup is a massive benefit; I'd guess that matters more than haploid selection. Also, low-fidelity non-destructive DNA sensing would allow one to select at the haploid level quickly and without culturing the haploids; see the appendix "Cheap DNA segment sensing".)

Another possible obstacle is that induced meiosis might be error-prone. Natural gametogenesis likely involves upregulating DNA protection and repair mechanisms, and involves culling meiotic cells with damaged DNA. In vitro meiosis might lack some of those mechanisms by default, so it might produce haploids with more DNA damage. Damaged cells can be filtered out iteratively, but this adds sand in the gears.

As I said above, with IMS the question of selection power is less "how much selection power", and more "how expensive for how much selection power". See the appendix "The costs of iterated meiotic selection" for more detail.

Auxiliary: Enhancements to meiotic methods

There are two hypothetical methods that would specifically enhance the power of methods that involve meiosis, such as iterated meiotic selection or IVG with gamete selection. These methods don't matter too much, as the bottlenecks are elsewhere, but I think they're interesting conceptually.

[145] Guglielmo, Maria Cristina, and David F. Albertini. 'The Structural Basis for Coordinating Oogenesis and Folliculogenesis'. In Oogenesis, edited by Giovanni Coticchio, David F. Albertini, and Lucia De Santis, 63–73. London: Springer, 2013. https://doi.org/10.1007/978-0-85729-826-3_5.

[146] Sagi, Ido, Gloryn Chia, Tamar Golan-Lev, Mordecai Peretz, Uri Weissbein, Lina Sui, Mark V. Sauer, Ofra Yanuka, Dieter Egli, and Nissim Benvenisty. 'Derivation and Differentiation of Haploid Human Embryonic Stem Cells'. Nature 532, no. 7597 (April 2016): 107–11. https://doi.org/10.1038/nature17408.

The first is *hypercrossover*. You somehow influence the diploid cell as it meioses, to make it have several crossovers per chromosome rather than the typical 1—3.

Segments would move around more freely, thus making themselves available to be selected-for more independently and therefore rapidly, e.g. in iterated meiotic selection.

Hypercrossover would also, in theory, increase the ceiling for single-round meiosis. If you made very many gametes (billions?), you'd do noticeably better with hypercrossover than without. (Maybe with only thousands; someone could check.)

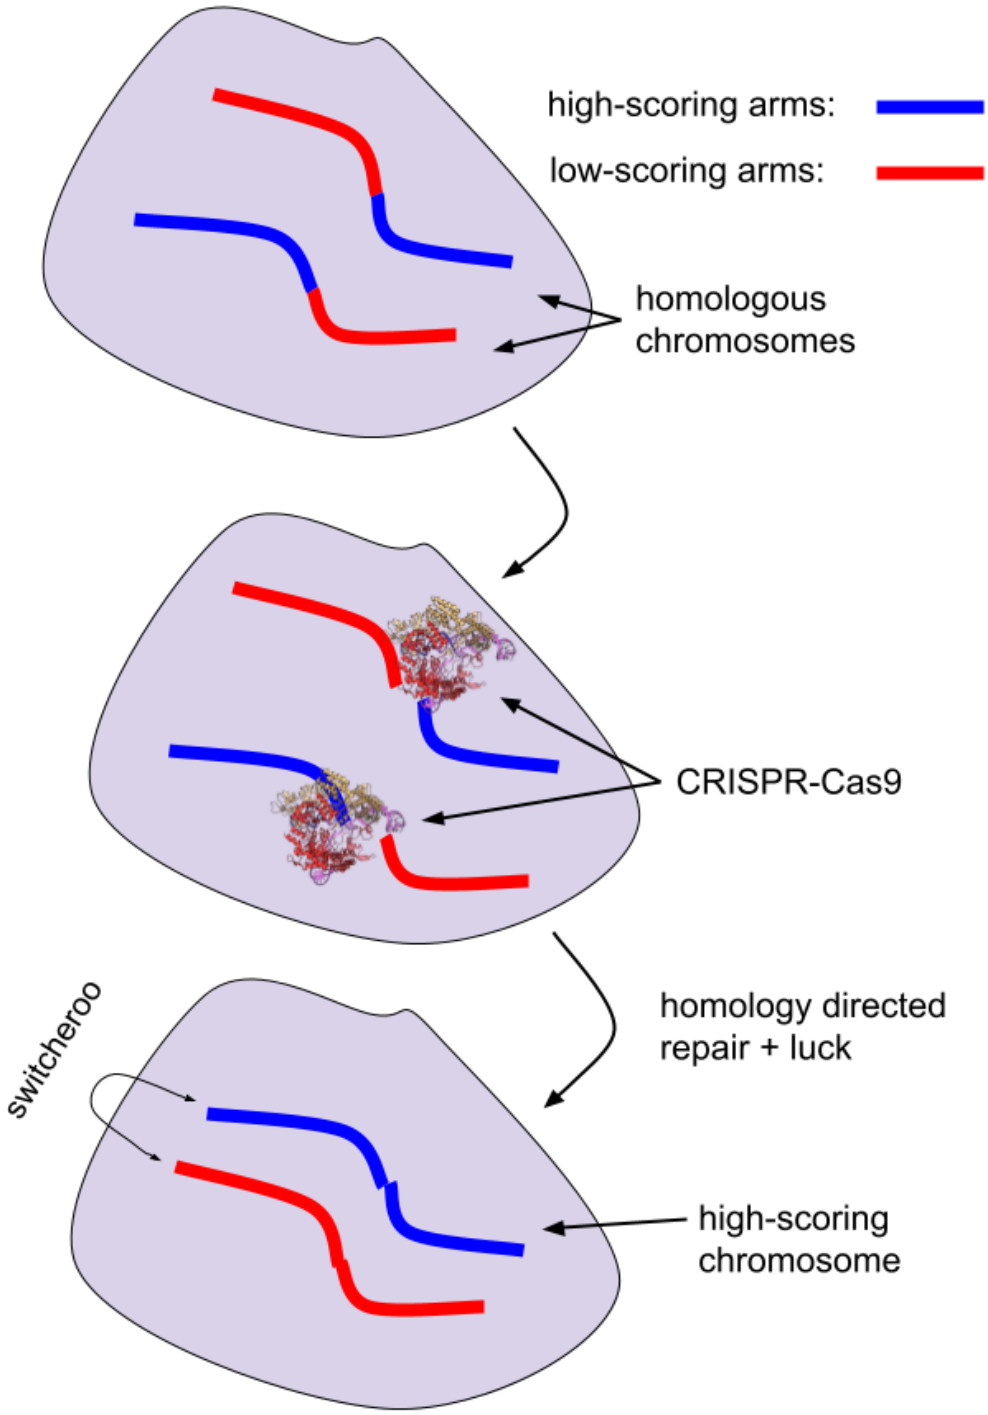
Finally, hypercrossover would be genuinely powerful combined with a round of chromosome selection. The highest-scoring chromosome would be higher-scoring by some multiple, like 2x or more, compared to the highest-scoring chromosome with only 1 crossover. (As a touchpoint, consider that with 3 crossovers, you're doing something like 4-segment selection, which is $\sqrt{4} = 2$ times as powerful as simple chromosome selection.)

Crossover rates can be increased in mice by knocking out a crossover-limiting gene[147]. I think the increase is about 30% here, based on the centimorgan numbers. This method would be bad for several reasons, and in general hypercrossover might cause DNA damage; but gentler methods, combined, could hypothetically produce several crossovers per chromosome.

The second auxiliary method is *targeted crossover*. You somehow influence the location of crossovers. If the influence is strong, then e.g. IMS becomes much cheaper, as you don't need to produce nearly as many haploids waiting for your desired crossover. If you can influence all the chromosomes at once in an in vitro gametogenesis protocol, you could have a vastly cheaper implementation of optimal gamete selection.

I don't know if targeted crossover is plausibly feasible. A possible method would be targeted double-strand breaks in diploid cells. You have two homologous chromosomes C and C', e.g. the two chromosomes of index 1. You have a desired crossover point. So you cleave both chromosomes at the crossover point using CRISPR-Cas9. Then you sit and hope that the cell's homology directed repair mechanisms swap the homologous arms, effectively making a crossover. (You'd still need some way of getting the resulting recombined chromosome into a cell to use as a gamete or zygote, e.g. via chromosome selection.) An illustration:

[147]Tsui, Vanessa, Ruqian Lyu, Stevan Novakovic, Jessica M. Stringer, Jessica E.M. Dunleavy, Elissah Granger, Tim Semple, et al. 'Fancm Has Dual Roles in the Limiting of Meiotic Crossovers and Germ Cell Maintenance in Mammals'. Cell Genomics 3, no. 8 (August 2023): 100349. https://doi.org/10.1016/j.xgen.2023.100349.



(The protein pictured is Cas9, visualization from Synpath.)

Method: Iterated multiplex CRISPR editing

See GeneSmith and kman (2025) for the source of the estimates of the strength of editing[148]. These estimates haven't been peer reviewed. They are based on a PGS that correlates at 0.35 with IQ, so they aren't directly comparable to estimates for selection methods given in this article, which assume 0.4.

Method:

1. Apply a set of CRISPR editors (e.g. base or prime editors) to some stem cells.
2. Culture the cells, perhaps monoclonally; if possible, sequence them to check for on-target edits made and off-target damage done.
3. Repeat 1 and 2, correcting errors if possible.
4. Make a baby from the edited cell.

(Aside: It would also be possible to safely edit an embryo directly. You'd use prime or base editors to edit

[148] GeneSmith, and kman. 'How to Make Superbabies', 19 February 2025. https://www.lesswrong.com/posts/DfrSZaf3JC8vJdbZL/how-to-make-superbabies.

the zygote (1-cell embryo), and you'd apply some treatment that deactivates the editors after a couple days. Later you sequence the embryo to ensure that there isn't any dangerous DNA damage. Since you edit at the 1-cell stage and then deactivate the editors, there shouldn't be mosaicism, so you can be fairly confident that the lack of damage in cells you biopsy and sequence really does mean that the embryo itself is fine. However, the restriction to only edit at the 1-cell stage means you can only do 1 round (perhaps multiplexed, like 10 or 20 edits). So this method isn't a strong GV method.)

DNA damage from editing

A general issue is dealing with DNA damage caused by editors. There are at least four kinds of damage:

- *Chromosome translocation/fragmentation.* Some editors, such as the original Cas9 system, cause double-strand breaks in the DNA. Usually double-strand breaks are repaired by reattaching the two ends of the chromosome that was broken. But sometimes the ends aren't repaired, or are repaired by stitching together two non-homologous broken chromosomes.
- *Off-target edits.* All CRISPR editors have some amount of mistargeting, where they simply edit some locus other than the intended target. The guide RNA that makes the editor target some DNA segment can also sometimes bind to other similar DNA segments at different sites in the genome.
- *Indels.* At the editing site, when an editor breaks/removes/unzips some DNA, unwanted nucleotides sometimes attach to the exposed DNA ends, causing unwanted insertions. Editors can also cause unwanted deletions.
- *Bystander edits.* Base editors cause bystander edits. After binding to the target site, they convert the target base to another nucleotide. But they sometimes also convert nearby bases of the same type as the target.

Depending on the rates of errors, this damage can probably be greatly attenuated by constantly sequencing the cells and discarding ones with any / many errors. They can also potentially be fixed by editing again (though this adds complexity and potentially significant delay, as you'd have to reengineer a new guide RNA to make the unexpected edit). Base and prime editors have much lower error rates than older double-strand breaking CRISPR systems, but still cause damage. DNA damage from editors might be especially bad compared to natural mutation, given that it introduces probably very rare de novo mutations.

Naive ESC editing

As discussed above in "Using natural reproductive DNA", it may be possible to maintain an ESC in a naive state and edit it, and then use that cell as a zygote. There will likely be issues with maintaining the epigenomic state, especially the imprinting. As with any method that tries to operate with naive ESCs to be used directly as zygotes, this might be a progressive problem: we might see significant early achievable gains, but then pushing forward another couple iterations (requiring several passages and proliferation) might continually degrade the epigenomic state.

Hulk sperm

Another EC bypass method, named by GeneSmith, is as follows.

1. Biopsy the basement membrane of a testicle to obtain gonadal tissue, including spermatogonia.
2. Edit the spermatogonia in vitro get target edits.
3. Make a monoclone and verify the integrity of the editing spermatogonia.
4. Edit in a fluorescent reporter near a gene that's only expressed in sperm.
5. Re-transplant the edited spermatogonia back into the testicle.
6. Wait a few months.
7. Obtain a sperm sample.
8. Isolate the magic glowing green sperm, e.g. using FACS.

One minor issue is that, if I understand correctly, the fluorescent reporter is a whole gene knock-in, so it would be inserted via normal double-strand-break CRISPR-Cas9 editing. Double-strand-break editing is prone to DNA damage such as indels around the break site. One always verifies the embryos at the final stage; so hopefully this just means that some smallish fraction of embryos will be unsafe to implant.

A larger issue is maintaining the spermatogonia in their support niche. If they fall out of their niche, they might not be able to proliferate in general, or differentiate into sperm, or have all the right epigenomic imprints.

An alternative method would use iPSCs instead of extracted spermatogonia. Then, after editing, you'd convert the iPSCs to spermatogonial stem cells, and then transplant them.

As mentioned in the earlier section "Features of GV and EC methods that affect compatibility", an issue with spermatogonial transplant into testicles is that GV (e.g. edited) spermatogonia are likely to be stressed and imperfectly maintained, and therefore are likely to be outcompeted by normal spermatogonia and die out. A possible workaround is to just transplant a large amount of GVed SSCs into the testes. I think that FACS can be quite sensitive, so if even a few glowing sperm make it out, they may be retrievable.

Magic rainbow sperm

Here's a weirder kind of editing: in vivo editing. This is known as the "we zap your balls until you make magic rainbow sperm", which is not a good method. How it works:

1. Transfect a CRISPR editing package into testes in vivo. The package contains instructions for several desired target edits (that is, several guide RNAs), and also a fluorescent reporter protein to be expressed only in sperm.
2. Zap the testicles to electroporate lots of cells, including hopefully many spermatogonia.
3. Repeat 1 and 2 using different colors for each round, so you'll be able to tell which ones got edits from many different rounds.
4. Wait a few months.
5. Obtain a sperm sample.
6. Isolate the magic glowing rainbow sperm, e.g. using FACS to find the brightest and most colorful sperm.

This isn't the best method: the edited spermatogonia would be stressed, knocking in several fluorescent proteins might be an issue, the editing would be inefficient because it wouldn't reach many cells, and surrounding tissue might be damaged by the editing and the electroporation. But it is the funniest method.

Conclusion

See the Takeaways. :)

Some other random small things you could do that might be fun:

- Come up with more feasible ways of doing GV and EC.
- Make a dashboard for combining GV methods in complex ways and showing the resulting GV power.
- Make an IMS simulator / game. Use AI or something to come up with clever ways of scheduling meiosis, sequencing, culturing, and cell fusion, to get large effects at low cost. Try adding in hypercrossover, targeted recombination, or chromosome transplantation.
- Review / replicate / fact check this article.
- Compute some GV power numbers for making multiple children for one couple, taking into account the goal of having children with genetic overlap that's not much more than the overlap of siblings.
- Simulate hypercrossover. E.g., in practice how does the distribution of haploid meiotic daughter genomes change with different numbers of crossovers? How does the in-practice maximum chromosome score change?

My work is not being supported and I am approximately out of money. If you want me to continue working on what I think I should be working on, you can support my work by donating here. (I'll edit this message if I receive enough to last 1.5 years.)

Appendix: In vitro spermatogenesis studies

Here are some of the main studies:

- Some early studies include Tesarik et al. (1999)[149], Cremades et al. (1999)[150], and Tanaka et al. (2003)[151]. Notably, Tesarik derived two apparently healthy live human births from round

[149]Tesarik, Jan, Mustafa Bahceci, Cenk Özcan, Ermanno Greco, and Carmen Mendoza. 'Restoration of Fertility by In-Vitro Spermatogenesis'. The Lancet 353, no. 9152 (13 February 1999): 555–56. https://doi.org/10.1016/S0140-6736(98)04784-9.

[150]Cremades, N., R. Bernabeu, A. Barros, and M. Sousa. 'In-Vitro Maturation of Round Spermatids Using Co-Culture on Vero Cells'. Human Reproduction (Oxford, England) 14, no. 5 (May 1999): 1287–93. https://doi.org/10.1093/humrep/14.5.1287.

[151]Tanaka, Atsushi, Motoki Nagayoshi, Shoichiro Awata, Yoshifumi Mawatari, Izumi Tanaka, and Hiroshi Kusunoki. 'Completion of Meiosis in Human Primary Spermatocytes through in Vitro Coculture with Vero Cells'. Fertility and Sterility 79 (1

spermatids produced from primary spermatocytes in vitro. However, another embryo implantation attempt resulted in an ectopic pregnancy, which is really bad—the embryo starts growing somewhere other than the uterus. Tanaka et al. got ok efficiency: around 5% spermatids from primary spermatocytes, on a culture of Vero cells (iPSCs from an African green monkey). These early studies had poor verification of haploid genome content, and very poor measurement of the epigenomic state of their in vitro-derived haploids. Tanaka et al. (2009) continued applying their Vero cell culture method, and were able to mature round spermatids into elongated spermatids and even some spermatozoa[152].

- Easley et al. (2012) cultured human iPSCs with mouse SSCs (spermatogonial stem cells). They found that about 40% became spermatogonia-like (according to a few markers), and that about 3.9% of them became haploid, some with other round spermatid-like characteristics. However, they didn't do scRNA-seq or methylation sequencing (because at the time, that technology hadn't yet been developed), so we don't know what their haploids look like epigenomically; they checked just 2 imprinting sites[153].
- Perrard et al. (2016) used a bioreactor to culture adult human seminiferous tubule tissue. SSCs eventually produced some spermatids and some mature spermatozoa. They didn't evaluate the epigenomic state of their spermatozoon-like cells[154].
- de Michele et al. (2018) took fragments of testes from prepubertal children who had died and cultured them for several months, finding that after a couple weeks some of the SSCs in the starting tissue had produced round spermatid-like cells, though again they were not well-characterized epigenomically[155].
- Sun et al. (2018) cultured human SSCs with human Sertoli cells filtered from adult testicular tissue, in a gel to make 3D tissue. They say the SSCs differentiated into round spermatids at a rate of up to 17.9%. They also did RNA sequencing and bisulfite methylation sequencing, showing significant similarity between their round spermatids and natural ones, though the similarity was far from perfect and I don't understand the results[156]. (They also showed that their *human* round spermatids could fertilize *mouse* oocytes.)
- Yuan et al. (2020) cultured human fetal gonadal ridge tissue, which probably contained prospermatogonia, and found that after many weeks there were around 5% round spermatid-like cells. In a first, they used single-cell whole genome sequencing to confirm that these haploids had proper DNA content from meiosis (the right number of chromosomes without major gain or loss). The spermatid-like cells could fertilize oocytes, though only one of 16 made it to the blastocyst stage (about a week into development). They only checked a couple imprinting regions, so we don't know to what extent epigenetic correction occurred[157].
- Hwang et al. (2020) turned human PGCLCs into prospermatogonia-like cells by culturing them with mouse testicular cells. They evaluated their cells with scRNA-seq, compared to natural prospermatogonia of a comparable stage, and found that they were broadly similar but with notable differences. I don't know how to evaluate the differences, but they do note that their prospermatogonia-like cells are enriched for some gene expression that's associated with infertility[158].

March 2003): 795–801. https://doi.org/10.1016/S0015-0282(02)04833-1.

[152]Tanaka, Atsushi, Motoi Nagayoshi, Shoichiro Awata, Izumi Tanaka, and Hiroshi Kusunoki. 'Differentiation of Human Round Spermatids into Motile Spermatozoa through in Vitro Coculture with Vero Cells'. Reproductive Medicine and Biology 8, no. 4 (21 August 2009): 169–75. https://doi.org/10.1007/s12522-009-0030-0.

[153]Easley, Charles A., Bart T. Phillips, Megan M. McGuire, Jennifer M. Barringer, Hanna Valli, Brian P. Hermann, Calvin R. Simerly, et al. 'Direct Differentiation of Human Pluripotent Stem Cells into Haploid Spermatogenic Cells'. Cell Reports 2, no. 3 (27 September 2012): 440–46. https://doi.org/10.1016/j.celrep.2012.07.015.

[154]Perrard, Marie-Hélène, Nicolas Sereni, Caroline Schluth-Bolard, Antonine Blondet, Sandrine Giscard d Estaing, Ingrid Plotton, Nicolas Morel-Journel, Hervé Lejeune, Laurent David, and Philippe Durand. 'Complete Human and Rat Ex Vivo Spermatogenesis from Fresh or Frozen Testicular Tissue1'. Biology of Reproduction 95, no. 4 (1 October 2016): 89, 1–10. https://doi.org/10.1095/biolreprod.116.142802.

[155]Michele, Francesca de, Jonathan Poels, Maxime Vermeulen, Jérôme Ambroise, Damien Gruson, Yves Guiot, and Christine Wyns. 'Haploid Germ Cells Generated in Organotypic Culture of Testicular Tissue From Prepubertal Boys'. Frontiers in Physiology 9 (9 October 2018). https://doi.org/10.3389/fphys.2018.01413.

[156]Sun, Min, Qingqing Yuan, Minghui Niu, Hong Wang, Liping Wen, Chencheng Yao, Jingmei Hou, et al. 'Efficient Generation of Functional Haploid Spermatids from Human Germline Stem Cells by Three-Dimensional-Induced System'. Cell Death & Differentiation 25, no. 4 (April 2018): 749–66. https://doi.org/10.1038/s41418-017-0015-1.

[157]Yuan, Yan, Laihua Li, Qing Cheng, Feiyang Diao, Qiao Zeng, Xiaoyu Yang, Yibo Wu, et al. 'In Vitro Testicular Organogenesis from Human Fetal Gonads Produces Fertilization-Competent Spermatids'. Cell Research 30, no. 3 (March 2020): 244–55. https://doi.org/10.1038/s41422-020-0283-z.

[158]Hwang, Young Sun, Shinnosuke Suzuki, Yasunari Seita, Jumpei Ito, Yuka Sakata, Hirofumi Aso, Kei Sato, Brian P. Hermann, and Kotaro Sasaki. 'Reconstitution of Prospermatogonial Specification in Vitro from Human Induced Pluripotent Stem Cells'. Nature Communications 11, no. 1 (9 November 2020): 5656. https://doi.org/10.1038/s41467-020-19350-3.

Appendix: Cheap DNA segment sensing

The problem of cheap DNA sensing

Suppose you have some cells, such as sperm or embryos, and you want to find out roughly what their genomes are, very cheaply and maybe without destroying them. If that's all you have, you might be in trouble. It's hard to detect much about a genome without some sort of DNA analysis that's destructive and expensive (SNP array genotyping, DNA sequencing), and that requires the somewhat labor-expensive step of isolating each cell individually.

But suppose further that the cells you have are descended from a cell with a known genome. E.g. you have sperm from a father whose whole genome you've sequenced, or you have stem cells created by combining a few different people's cells. Now the task is much more doable. In a sense, you already mostly know the genome of the cell you're interrogating! You just need to find out which large segments from the parent cell(s) made it into this cell.

Basic chromosome sensing

As a toy example, suppose you take two cells from two different people, and then create a single cell by randomly picking a full complement of chromosomes from the two given cells. (For example, you electrofuse the cells, and then wait for the tetraploidy to spontaneously correct through mitosis.) So, some process randomly picks two of the four given chromosomes of index 1; two of the four chromosomes of index 2; and so on for each index up to 23. Now, how to figure out the genome of this new cell? You mostly already know it. You just need to figure out which 2 of the 4 chromosomes at each index made it into the cell.

You could do this by sequencing or genotyping, though that would be relatively slow and a little expensive, as well as destructive. But it's also overkill. All you really need is < 3 bits of information per chromosome. In general, this seems to me like it ought to be doable, though I don't know of a method that would definitely work.

A method that might work and be non-destructive and fast, is FISH (fluorescence in situ hybridization). FISH puts fluorescent molecules into cells. The molecules are bound to DNA probes, which are targeted to bind only to specific shortish DNA sequences. They hybridize with the DNA, sticking to it for a while. Then you can literally see which chromosomes were bound or not.

So what we'll do is, for the four chromosomes of index 1 that we started with, we'll analyze their sequences to find one (or maybe a couple) identifying variants. E.g. maybe one of the chromosomes has SNPs AB, another has $A'B$, another has AB', and the last has $A'B'$. If we label A with an orange fluorescent protein, and label B with a green one, then based on orange/not-orange and green/not-green, we can tell which chromosomes are in our cell! E.g. a cell that's just green is $A'B$.

Sensing crossovers

Now a more complicated case: meiotic child cells. Natural meiosis—dividing from a diploid (46 chromosome) cell into a haploid (23 chromosome) cell—occurs as part of oogenesis and spermatogenesis; artificial in vitro meiosis might soon become possible. Meiosis recombines DNA from each pair of two homologous chromosomes in the parent diploid cell, producing a single chromosome that is, for example, the first half of the first parent chromosome, followed by the second half of the second chromosome. In human gametogenesis, it's believed that there are generally very few crossovers per chromosome; on average, maybe 1.1 in sperm, and maybe 1.8 in oocytes (though with substantial variation between people and within one person's gametes)[159]. If you're looking at a gamete, and you know the parent's genome, all you need to know is roughly where the crossovers (probably 1—3) happened in each chromosome, and then you know almost all of the genome of that gamete.

How to detect crossover points? If you have some DNA probes, such as FISH, you can detect a few SNPs, spread across the chromosome, that differentiate between the possible parent chromosomes that the segment around that SNP could have come from. Say you have a sperm and you label 5 loci, dividing the chromosome into 6 segments; you label dad's paternal chromosome's SNPs orange, and dad's maternal chromosome's SNPs green. Then if you see a pattern like [green, green, green, orange, orange], you can be fairly confident that the first half is maternal, and the last third is paternal. (You don't know where the crossover in the

[159] Bell, Avery Davis, Curtis J. Mello, James Nemesh, Sara A. Brumbaugh, Alec Wysoker, and Steven A. McCarroll. 'Insights into Variation in Meiosis from 31,228 Human Sperm Genomes'. Nature 583, no. 7815 (July 2020): 259–64. https://doi.org/10.1038/s41586-020-2347-0.

4th of 6 segments occurred; and you don't know for sure that there weren't more crossovers, e.g. one at the end in the middle of the 6th segment, or two inside of a single segment.)

Basically, you just label the DNA segments that are your targets, i.e. the ones you're trying to concentrate together in one cell as much as possible. Then you pick the cell with the most labels, i.e. the brightest cell.

An obstacle with sperm

There is a major issue with this protocol, though: I don't know of a method for labeling sperm DNA in a sequence-specific way. Sperm DNA is tightly condensed through protamination, and is therefore mostly inaccessible. Some sperm DNA is only histoned rather than protaminated, so it's possible that it could be labeled. But the chromatin is very tightly wound, so labeling molecules would have trouble getting to touch most of the DNA.

It's also possible that one could intervene to mildly decrease the degree of compaction, and then label the DNA when it's more accessible. For example, various treatments such as dithiothreitol are used to decondense sperm chromatin[160]. This probably causes some DNA fragmentation, but not at a very high rate. (In theory this starts to have some nontrivial effect on the broad-sense epigenomic state of the DNA, and therefore makes the case a bit less airtight that the labeled sperm should be exactly functionally equivalent to normal sperm. But probably it's perfectly fine, assuming that essentially the only function of protamination is for compaction (in order to swim), rather than silencing (which is anyway accomplished by methylations and the lack of much cellular machinery in the sperm).)

But, this might be fairly easy, or might be quite difficult, as far as I know.

Another issue is that the sperm chromosomes might be stuck together, e.g. by proteins or by being geometrically entangled.

Cheap sensing at scale with FACS

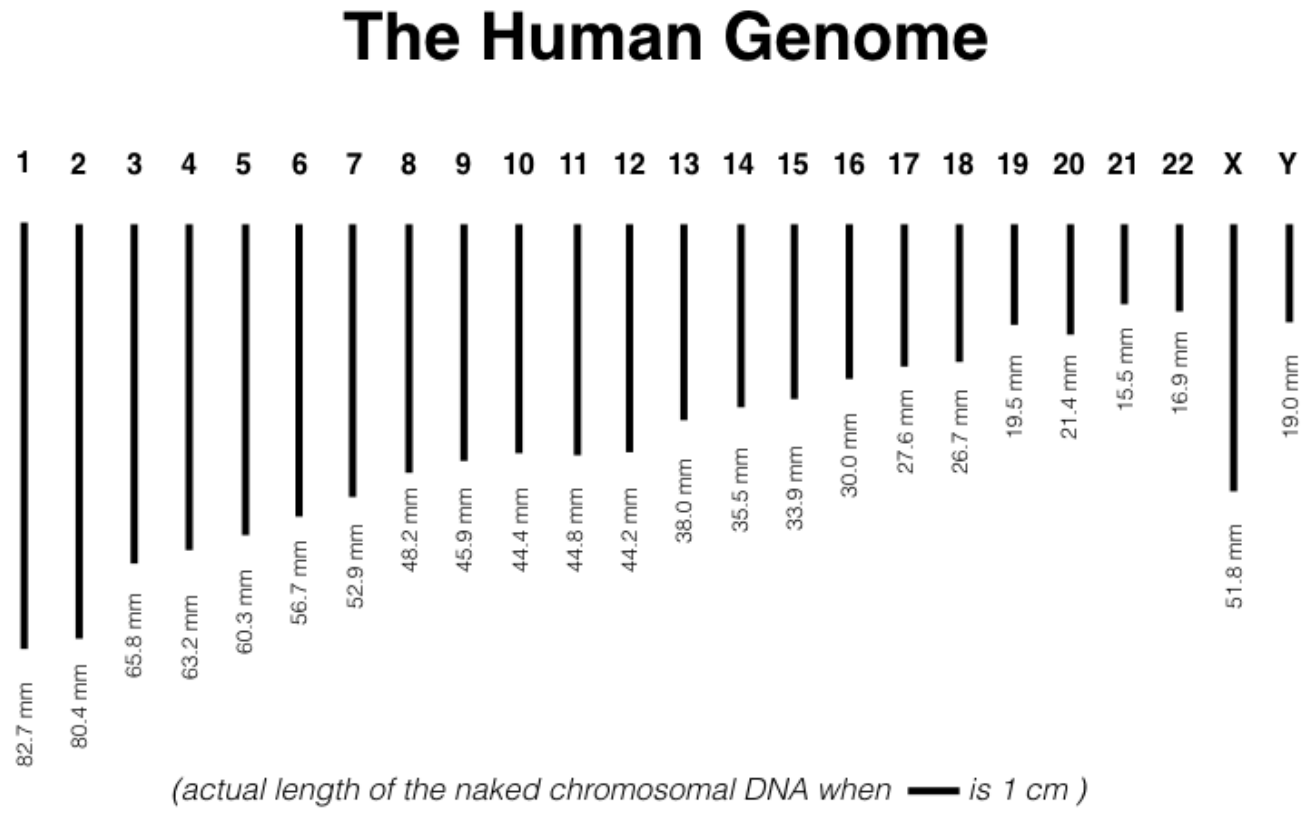
They key advantage of cheap DNA sensing would be that you can run assays at scale, e.g. to select from many thousands or millions of cells based on some DNA features. How do you select the brightest cell? Luckily there is existing technology that does something like this: fluorescence-activated cell sorting, or FACS, which sorts cells based on optical properties. I don't know enough to say confidently that one could sort in the necessary way, but my guess is you can. You could sort by brightness; or you could have multiple color labels and sort by activation of many colors. However, FACS might by default break up large DNA molecules.

Sensing chromosome index

In some contexts it may be possible to separate chromosomes by index without being able to tell anything else about the chromosomes.

One way to do this is simply to look with your eyes (classic!) through a microscope at the chromosomes, and see which ones look bigger or smaller. But look at the sizes of human chromosomes:

[160] Chankitisakul, Vibuntita, Nutthee Am-In, Theerawat Tharasanit, Tamas Somfai, Takashi Nagai, and Mongkol Techakumphu. 'Sperm Pretreatment with Dithiothreitol Increases Male Pronucleus Formation Rates After Intracytoplasmic Sperm Injection (ICSI) in Swamp Buffalo Oocytes'. Journal of Reproduction and Development 59, no. 1 (2013): 66–71. https://doi.org/10.1262/jrd.2012-104.



(Figure 1.3 from Gallegos (2022)[161].)

One might, for example, be able to visually distinguish chromosomes 1 and 2 from the other chromosomes. But good luck trying to eyeball the differences between 10, 11, and 12. If you can visually distinguish them well, you could even just guess. If you have two 2-way confusions (say, 21/22 and 1/2) and a 4-way confusion (9—12), you have to try about 16 times to get one correct one. Kinda expensive in embryos and labor but not that bad! (S. Eisenstat suggests: instead of guessing, do selection not to average chromosome score, but rather to average score of indistinguishable sets of chromosomes. E.g. you pick the highest-scoring donor to take all of chromosomes 9—12 from.) Though this may be infeasible because sperm chromosomes are quite compacted, so their length might not be visible.

Another possible method would be to lyse the cells and then centrifuge the sample, so that molecules separate by mass. Another hypothetical possibility would be pulsed-field gel electrophoresis. My gippities suggest perhaps Hoechst staining and flow cytometry. I don't know whether these or other methods would work, or would damage the DNA.

Appendix: Best crossover

Suppose you have a diploid cell, and you cause it to undergo meiosis. At index 1, the diploid cell's two homologous chromosomes will recombine, creating one or a couple crossover points. The meiotic child haploid cell will have a single chromosome at index 1, a cross between the two parent chromosomes.

Consider very many meiotic children created from one cell. The question is: what is the score of the top-scoring chromosome 1 out of all the crossed chromosome 1s?

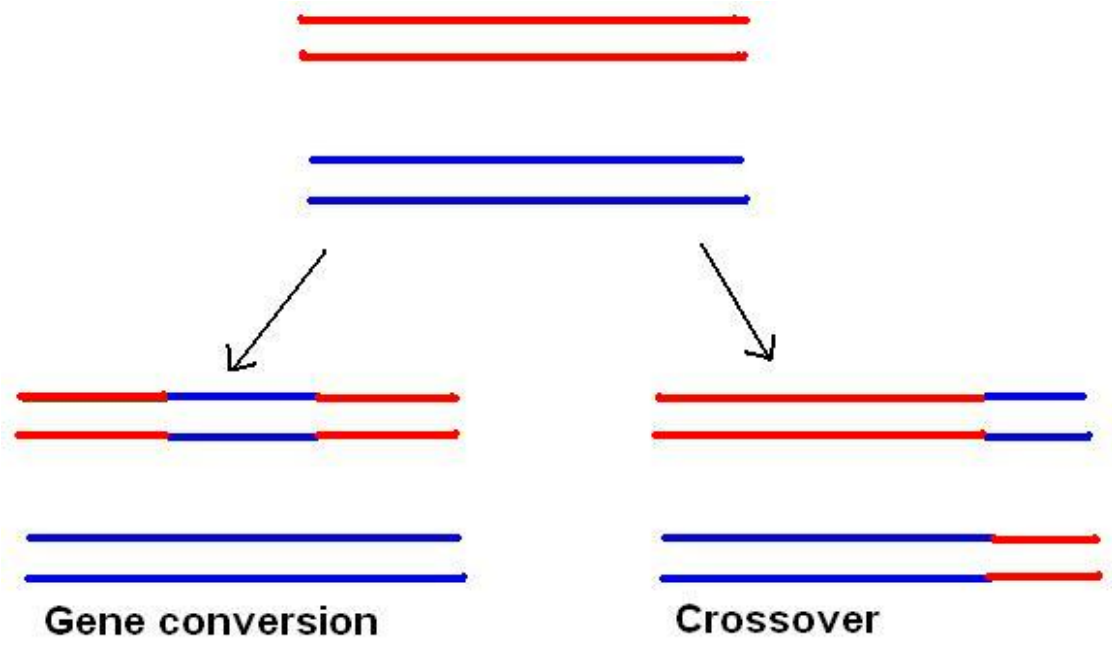
What does recombination do to DNA?

Backing up, what actually happens to DNA in the two chromosomes during meiotic recombination? The whole process of meiosis is complicated, but here we just need to know how DNA moves around between chromosomes.

There are two kinds of changes to DNA during meiosis: gene conversion and crossover. In gene conversion, a small segment of one chromosome is used as a template to copy over to the other homologous chromosome, overwriting that segment. In crossover, the two homologous chromosome totally "switch tracks", and they

[161] Gallegos, Maria. Fantastic Genes and Where to Find Them. Updated 2022-09-13. Accessed 16 February 2025. https://bookdown.org/maria_gallegos/where-are-genes/#preface.

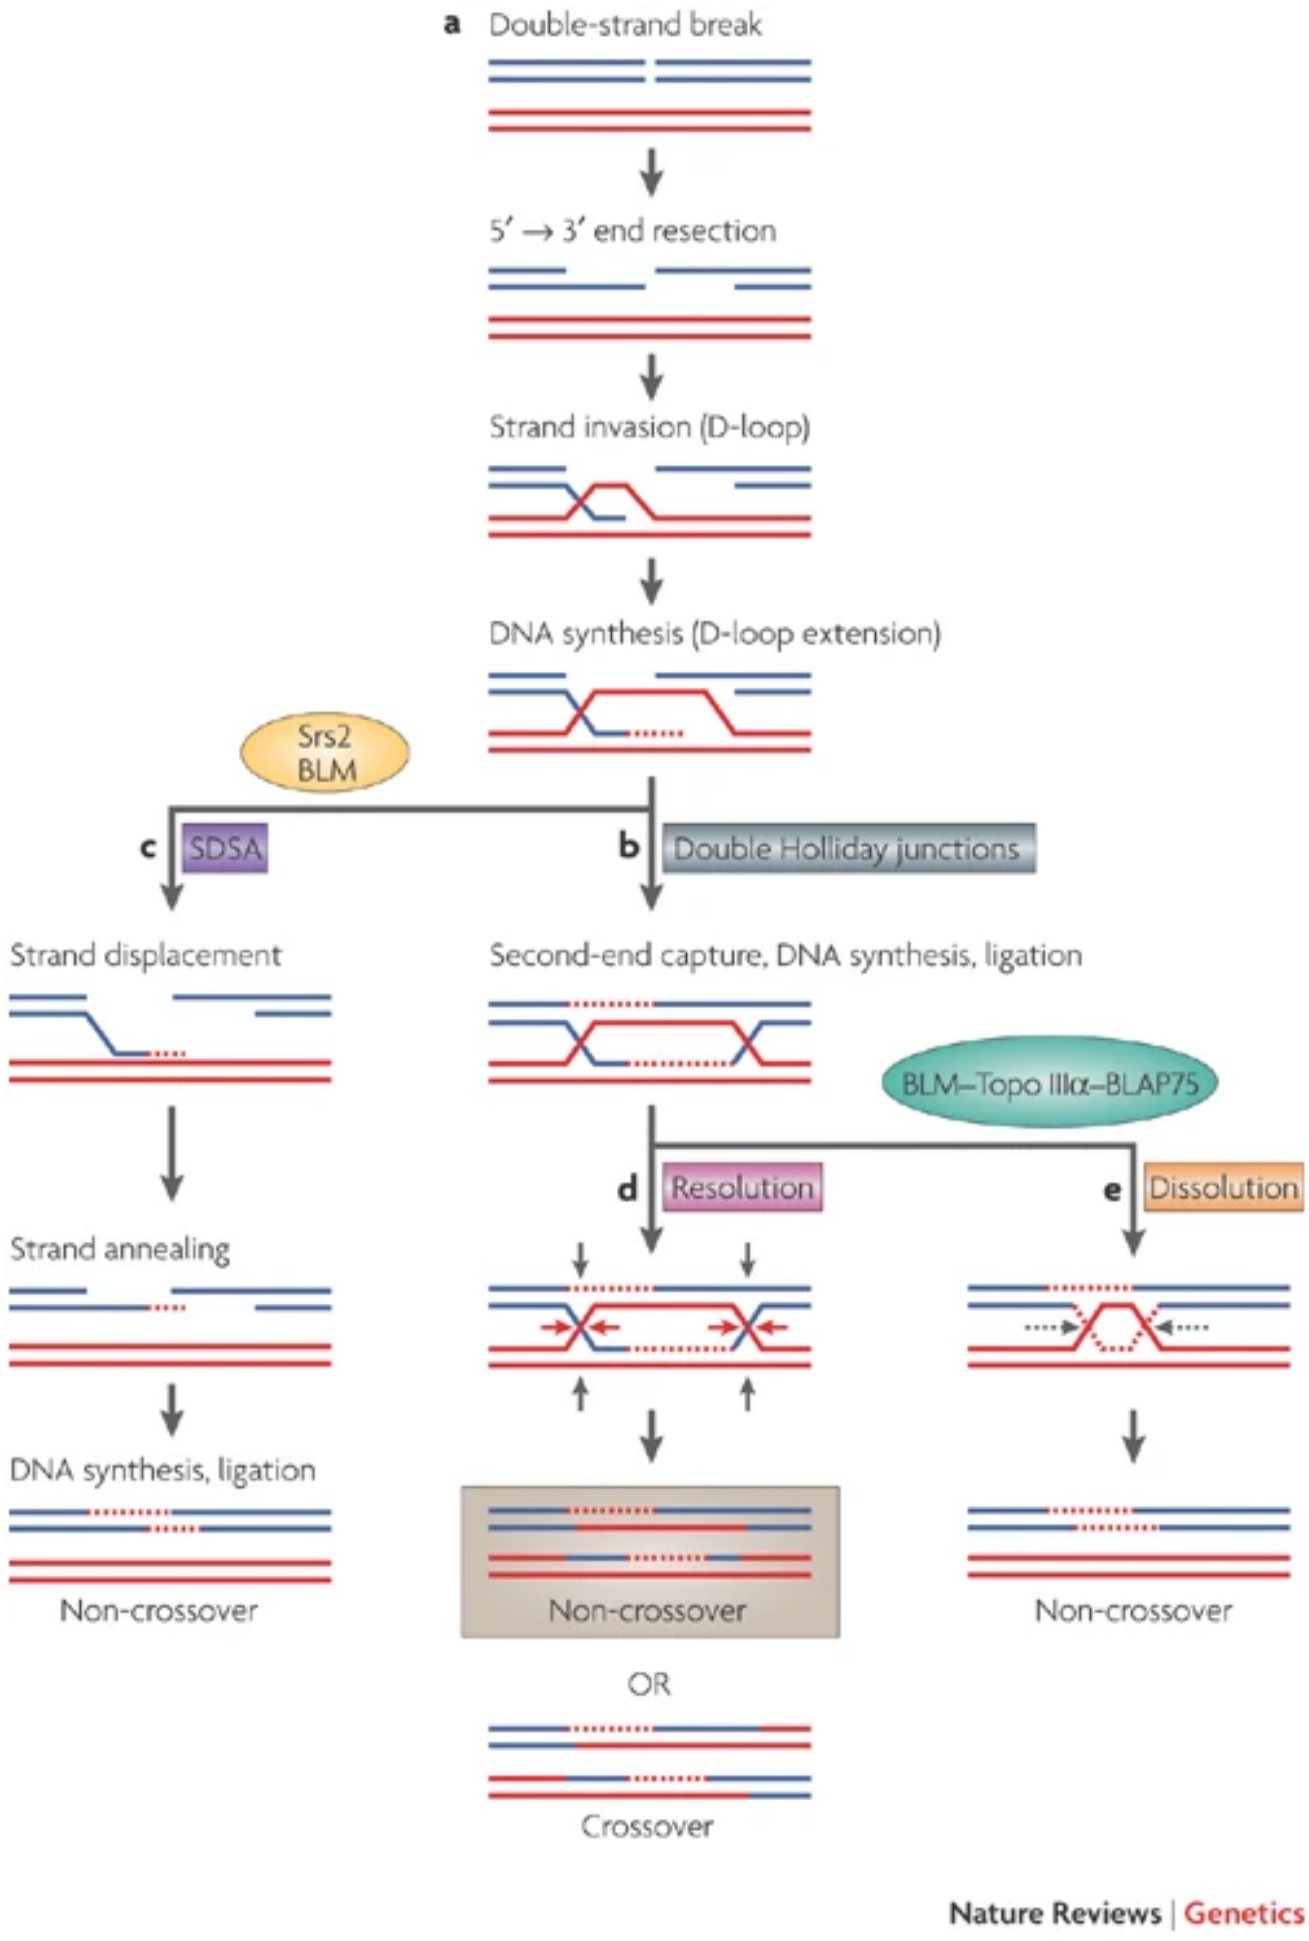
end up having exchanged some fraction of a whole arm, from one end of the chromosome up to the crossover point.



(Diagram from Mikael Häggström, Wikimedia: https://commons.wikimedia.org/wiki/File:Conversion_and_crossover.jpg)

In other words, a crossover of "A long time ago in a galaxy far, far away…" and "At the beginning of Elohim's creating of the skies and the land," could produce "At the beginning of Elohim's creating of a galaxy far, far away…".

How recombination actually happens is complicated:



(Figure 1 from Chen et al., 2007.[162])

[162]Chen, Jian-Min, David N. Cooper, Nadia Chuzhanova, Claude Férec, and George P. Patrinos. 'Gene Conversion: Mechanisms, Evolution and Human Disease'. Nature Reviews Genetics 8, no. 10 (October 2007): 762–75. https://doi.org/10.1038/nrg2193.

See also this beautiful animation by @WEHImovies (Drew Berry et al.), "DNA Break Repair by Homologous Recombination", which shows double Holliday junctions: https://www.youtube.com/watch?v=Xe-83tBcxhs

The main things we need to know:

- Gene conversion affects only a tiny portion of DNA, so we can largely ignore it.
- Sperm chromosomes have on average 1.1 crossovers or so, and oocytes are 1.8 or so, though both vary a lot between people and between gametes from the same person[163].

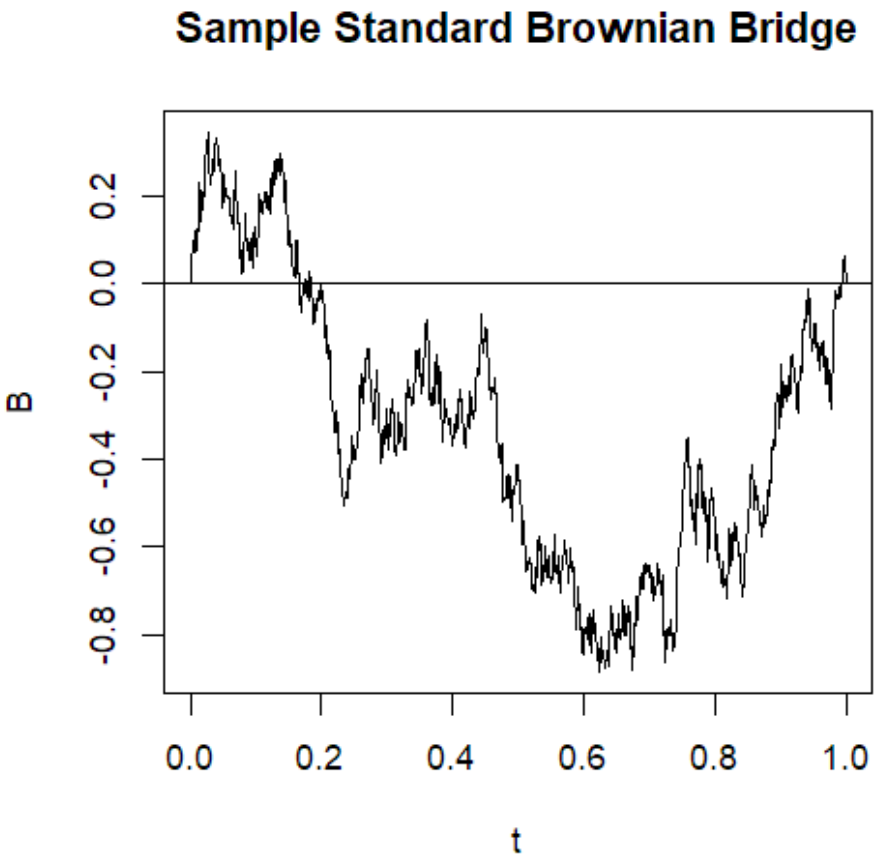
Math analysis of crossovers

As a modeling assumption, we can think of each chromosome as approximating a 1-D brownian motion starting at zero. The value of the brownian motion at, say, thirdway through the chromosome, is the score of the first third of the chromosome.

The difference between the local (e.g. SNP) scores of the two parent chromosomes also forms a brownian motion. As a simple base case, assume further that this brownian motion happens to end at 0. (This corresponds to the two chromosomes having exactly the same total score.)

Now the highest score among cross chromosomes will just be twice the maximum distance from 0 that this brownian bridge goes. (Plus the average score of the two chromosomes, which we can factor out of the analysis.) This problem has been treated by Kolmogorov and others; see e.g. "An excursion approach to maxima of the brownian bridge" for a source[164].

To illustrate what's happening, say we start with a brownian bridge:



(This diagram is from Felix Xiao[165], and the following ones are modifications of it.)

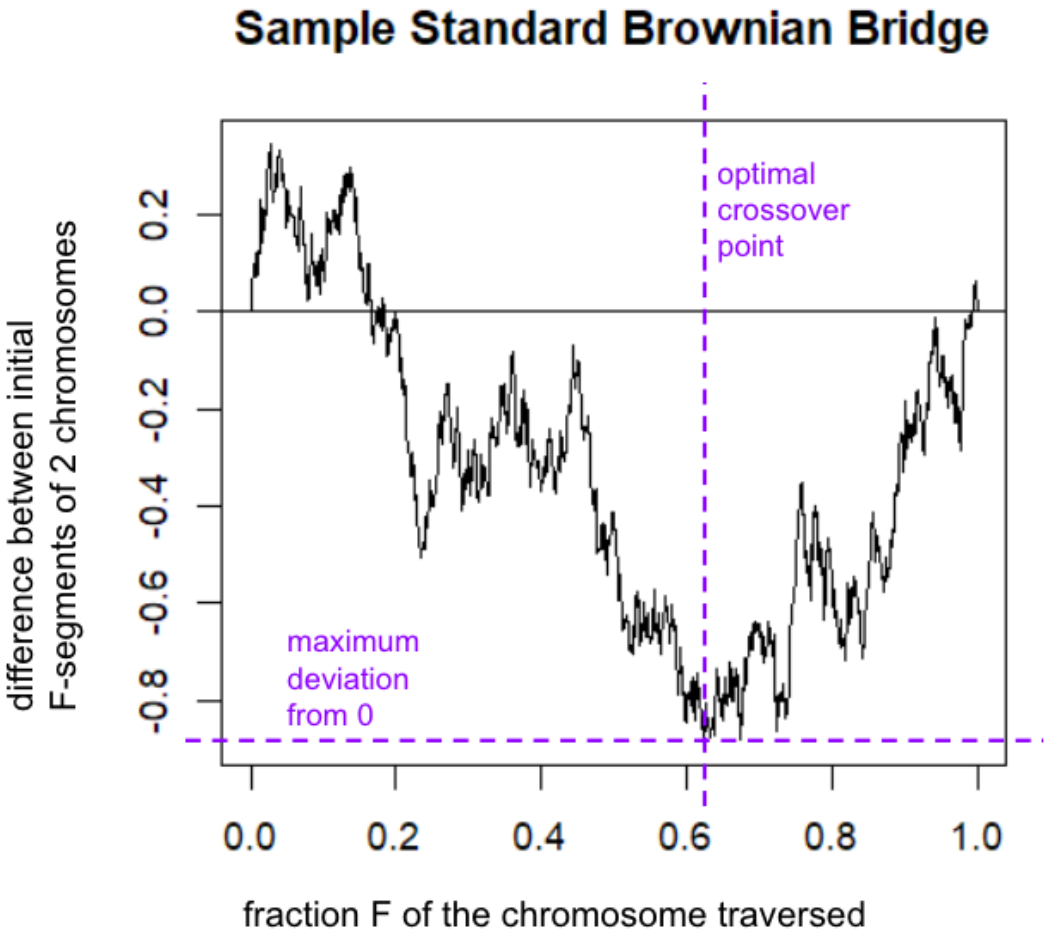
[163] Bell, Avery Davis, Curtis J. Mello, James Nemesh, Sara A. Brumbaugh, Alec Wysoker, and Steven A. McCarroll. 'Insights into Variation in Meiosis from 31,228 Human Sperm Genomes'. Nature 583, no. 7815 (July 2020): 259–64. https://doi.org/10.1038/s41586-020-2347-0.

[164] Perman, Mihael, and Jon A. Wellner. 'An Excursion Approach to Maxima of the Brownian Bridge'. Stochastic Processes and Their Applications 124, no. 9 (1 September 2014): 3106–20. https://doi.org/10.1016/j.spa.2014.04.008.

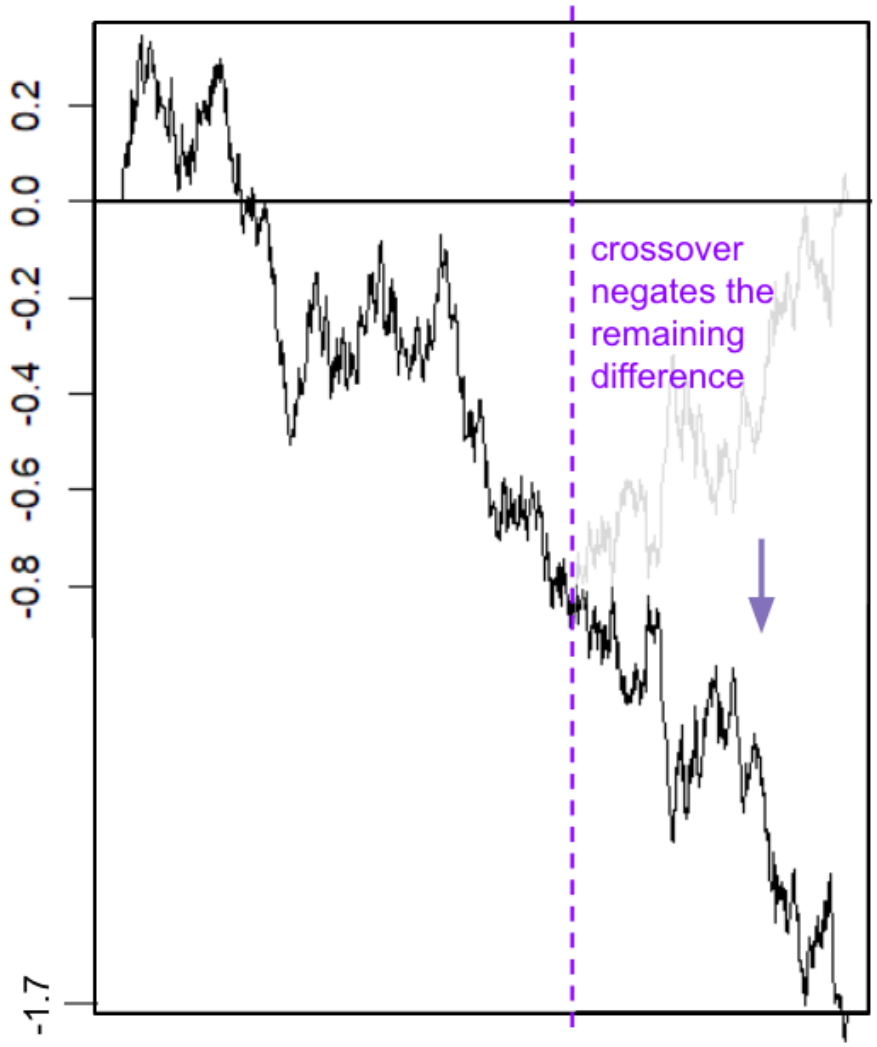
[165] Xiao, Felix. 'The Brownian Bridge Joint Max Position Distribution'. Undersampled - Felix Xiao, 24 January 2018. https://felixxiao.github.io/2018/01/brownian-bridge.

At position F along the x-axis, we view the height of the brownian bridge as being the cumulative differences of the scores of the two chromosomes, up to F along them. So e.g. if F is $1/2$, we see that one chromosome's first F, i.e. the first half, is about 0.7 lower-scoring than the first half of the other chromosome.

To find the best crossover point, we just look at the place where the bridge deviates the most from zero:



Then, if we get a crossover there, the resulting chromosome will have gotten the steepest difference in the first segment; followed by the steepest difference in the second segment, but with the sign inverted. ("Back in my day, we walked the whole length of the chromosome, uphill both ways.") The result looks like this:



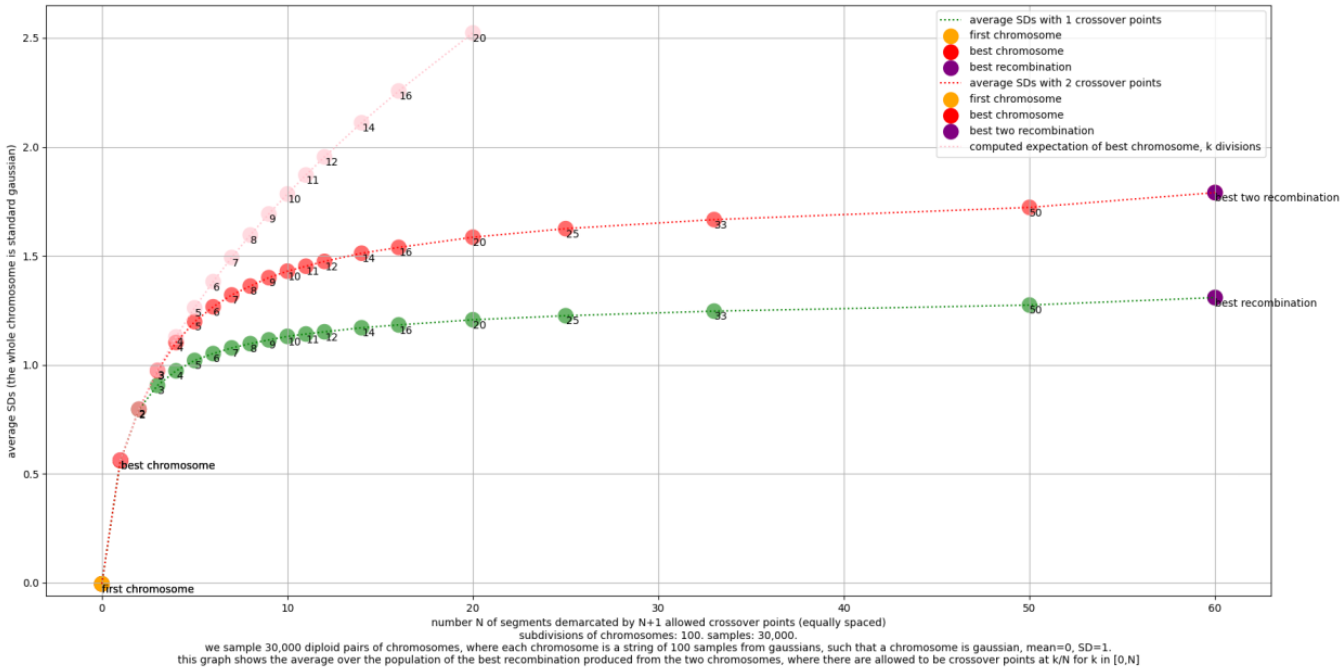
In this way, the crossed chromosome can have a score significantly far from zero, even if the parent chromosomes have the same average score.

For the full case, where the chromosomes can have different total scores, there are complications. E.g., we're uncertain about the endpoint of the bridge, and we're not just asking for the maximum distance to 0.

The case where there are 2 crossovers is additionally complicated.

Simulation results of crossovers

A simulation of recombination gives:



The simplest info in this graph is on the far right. It says that the best recombination with 1 crossover is about 1.3 SDs (in the scale of the whole chromosome), and the best recombination with 2 crossovers is about 1.8 SDs.

The rest of the graph shows the gain if you can only sense crossover locations with some precision. More precisely, the x-axis is the number of segments between allowed crossover points. At 1 on the x-axis, no crossover is allowed; at 2, only crossovers in the middle are allowed; at 3, crossovers at 1/3 and 2/3 are allowed; etc.

The pink line shows the results if you can freely mix and match the segments between chromosomes, however many crosses that takes.

We see that most of the benefit is gotten from 10 or even 5 segments. So sensing the chromosome of origin for 5 or 10 loci, spaced out through the chromosome, gets most of the benefit.

Here's the code for simulating random genomes: https://gist.github.com/tsvibt/89cec1b4fd7d54be04ba34a0059bcc4c. It computes maximum-score recombinations for each chromosome pair. By changing the number of segments per chromosome, you can simulate different levels of resolution used in recombinant chromosome selection. To replicate the numbers I give above for single crossover chromosomes, you can use something like this:

```
from shared_genome import *

def average_best_recombination(segment_count, sample_size):
    chromosome_pairs = [DiploidChromosome(segment_count=segment_count) for _ in range(sample_size)]
    return np.average([pair.max_recombination for pair in chromosome_pairs])

average_best_recombination(5, 10000)
```

E.g.:

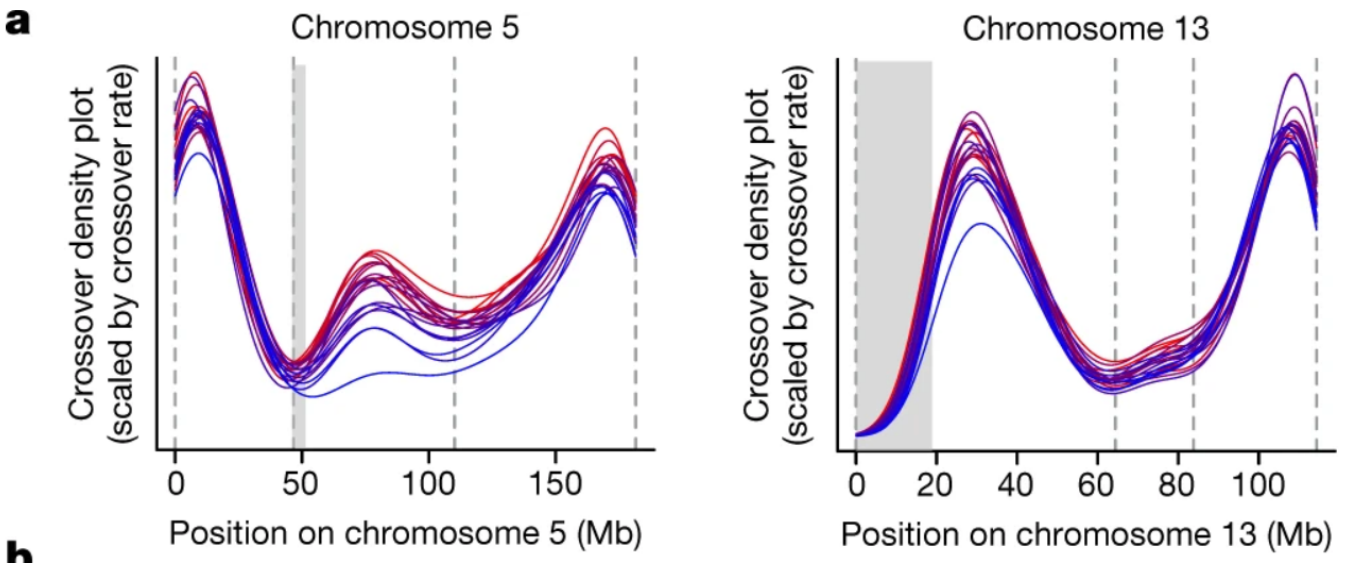
```
>>> average_best_recombination(5, 10000)
1.0188024501193658
>>> average_best_recombination(5, 10000)
1.012057646519421
>>> average_best_recombination(10, 10000)
1.1343196475916841
>>> average_best_recombination(10, 10000)
1.1329019005532246
>>> average_best_recombination(30, 10000)
1.2360039835782262
>>> average_best_recombination(30, 10000)
1.2496364605062995
```

Refinements to the crossover model

The above analysis diverges from reality in several ways. Some of them:

Non-uniform crossover points

Crossovers don't happen at uniformly random points along the chromosome. Rather, they're somewhat concentrated in some areas, especially distal regions:



(Figure 2a, from Bell et al., 2020[166].)

This pattern at least qualitatively matches the distribution of functional DNA in the chromosomes; e.g. the ends and very middle of a chromosome don't do much coding or regulation (but rather have functions related to the physical structure of the whole chromosome). So if we rescale by relevance to germline engineering, the non-uniformity of crossover location would be somewhat attenuated (I don't know how much).

Further, for many purposes, a 5x difference in commonality of a crossover at a given point is not so bad. For example, for recombinant chromosome selection, even if the target crossover point is 5x less common, you just need to look at 5x as many chromosomes. Since you only had to look at tens of chromosomes anyway to get pretty close to optimal, this is not a problem for sperm chromosome selection. For oocyte chromosome selection, it would be somewhat of a problem, as even having 30 eggs would usually be considered an abundance. Then again, for ESC chromosome selection, if there are enough epigenomically intact cells (e.g. via twinning or just growth), no loss of oocytes / embryos is needed in order to attempt chromosome selection.

Multiple crossovers

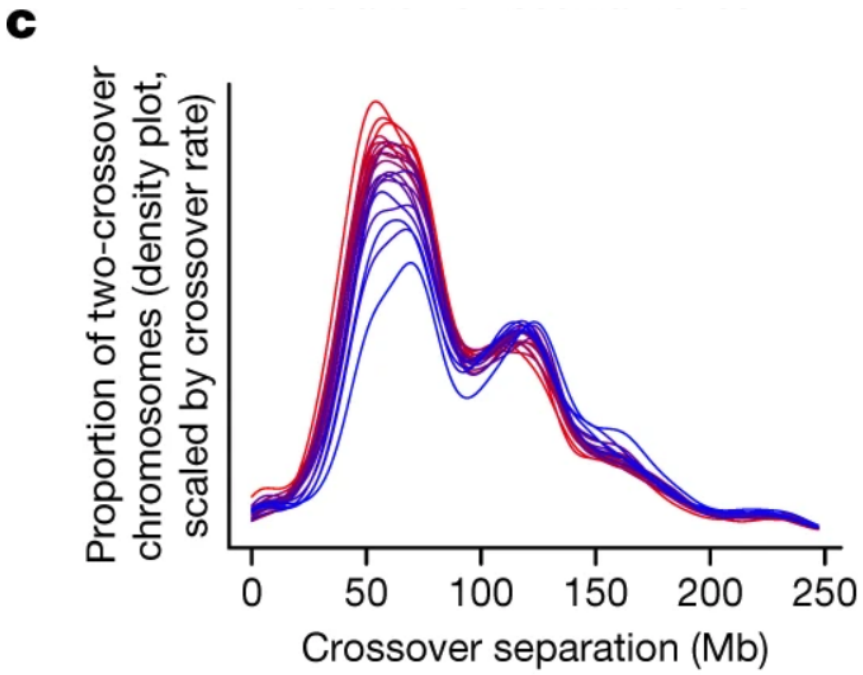
The average number of crossovers per chromosome in sperm differs a bit between men, from 1 to 1.2 or so. However, the number per sperm varies much more: the standard deviation in the total number of crossovers in a sperm is greater than 4. That means chromosomes with 2 or more crossovers must be fairly common[167].

In the simulations above, the maximum score for chromosomes with 1 crossover is about 1.3 chromosome-SDs. But the maximum score for chromosomes with 2 crossovers is significantly larger, at about 1.8 chromosome-SDs; with 5 or 6 distinguished segments, the top-scoring 1-crossover is around 1 SD, but the top-scoring 2-crossover is around 1.3 SDs. There would also be some gain from selecting 3-crossovers. So multiple crossovers might significantly increase the selection power.

On the other hand, the relative positions of multiple crossovers aren't random. Rather, they tend to be somewhat nearby:

[166]Bell, Avery Davis, Curtis J. Mello, James Nemesh, Sara A. Brumbaugh, Alec Wysoker, and Steven A. McCarroll. 'Insights into Variation in Meiosis from 31,228 Human Sperm Genomes'. Nature 583, no. 7815 (July 2020): 259–64. https://doi.org/10.1038/s41586-020-2347-0.

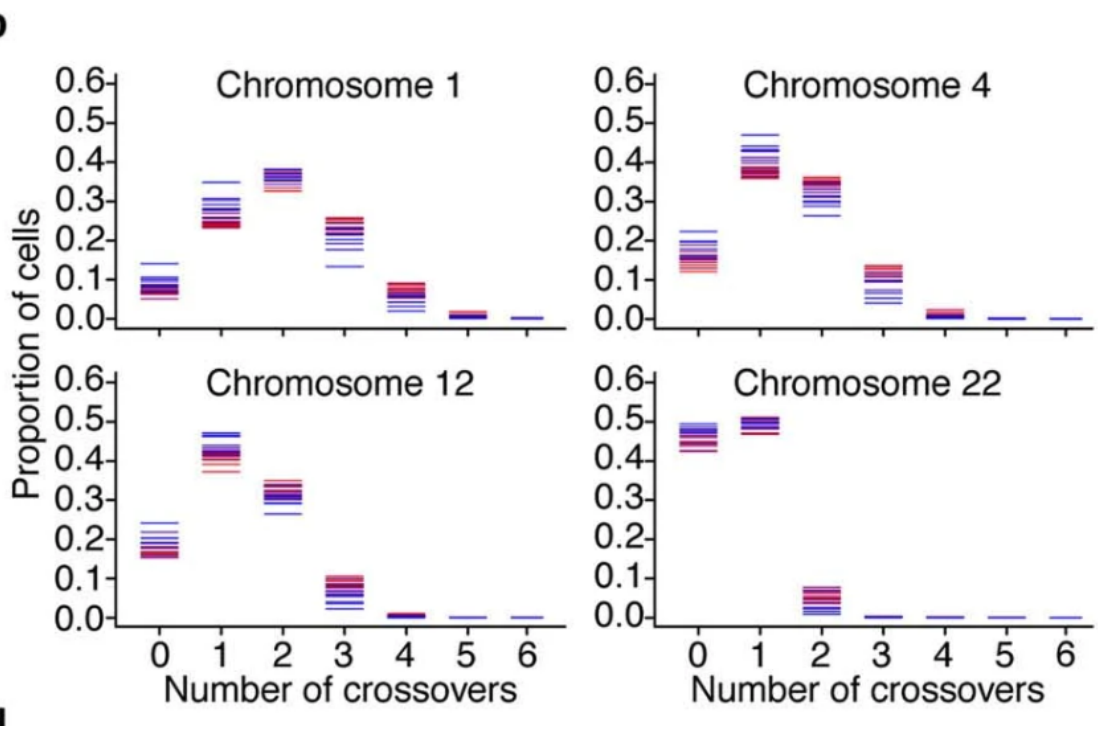
[167]Bell, Avery Davis, Curtis J. Mello, James Nemesh, Sara A. Brumbaugh, Alec Wysoker, and Steven A. McCarroll. 'Insights into Variation in Meiosis from 31,228 Human Sperm Genomes'. Nature 583, no. 7815 (July 2020): 259–64. https://doi.org/10.1038/s41586-020-2347-0.



(Figure 2c, from Bell et al., 2020.)

This somewhat attenuates the effect of selecting 2-crossovers, though probably not all that much.

Empirically there are many 3-crossover chromosomes in sperm, though the rates vary a lot between men and between chromosome index:



(Figure 8b, from Bell et al., 2020. The bars indicate different donors; the colors indicate high or low average crossover rate for that donor.)

Oocyte chromosomes have a much larger rate of crossovers than sperm—around 1.8 per chromosome, though again with a lot of variation.[168]

There is also non-crossover recombination, in the form of gene conversion. But this only transfers a tiny fraction of DNA between chromosomes—well under .01%[169]—and therefore doesn't much affect selection power.

[168] Ottolini, Christian S., Louise Newnham, Antonio Capalbo, Senthilkumar A. Natesan, Hrishikesh A. Joshi, Danilo Cimadomo, Darren K. Griffin, et al. '"Genome-Wide Recombination and Chromosome Segregation in Human Oocytes and Embryos Reveal Selection for Maternal Recombination Rates"'. Nature Genetics 47, no. 7 (July 2015): 727–35. https://doi.org/10.1038/ng.3306.

[169] Palamara, Pier Francesco, Laurent C. Francioli, Peter R. Wilton, Giulio Genovese, Alexander Gusev, Hilary K. Finucane, Sriram Sankararaman, et al. 'Leveraging Distant Relatedness to Quantify Human Mutation and Gene-Conversion Rates'. The American Journal of Human Genetics 97, no. 6 (December 2015): 775–89. https://doi.org/10.1016/j.ajhg.2015.10.006.

Appendix: The costs of iterated meiotic selection

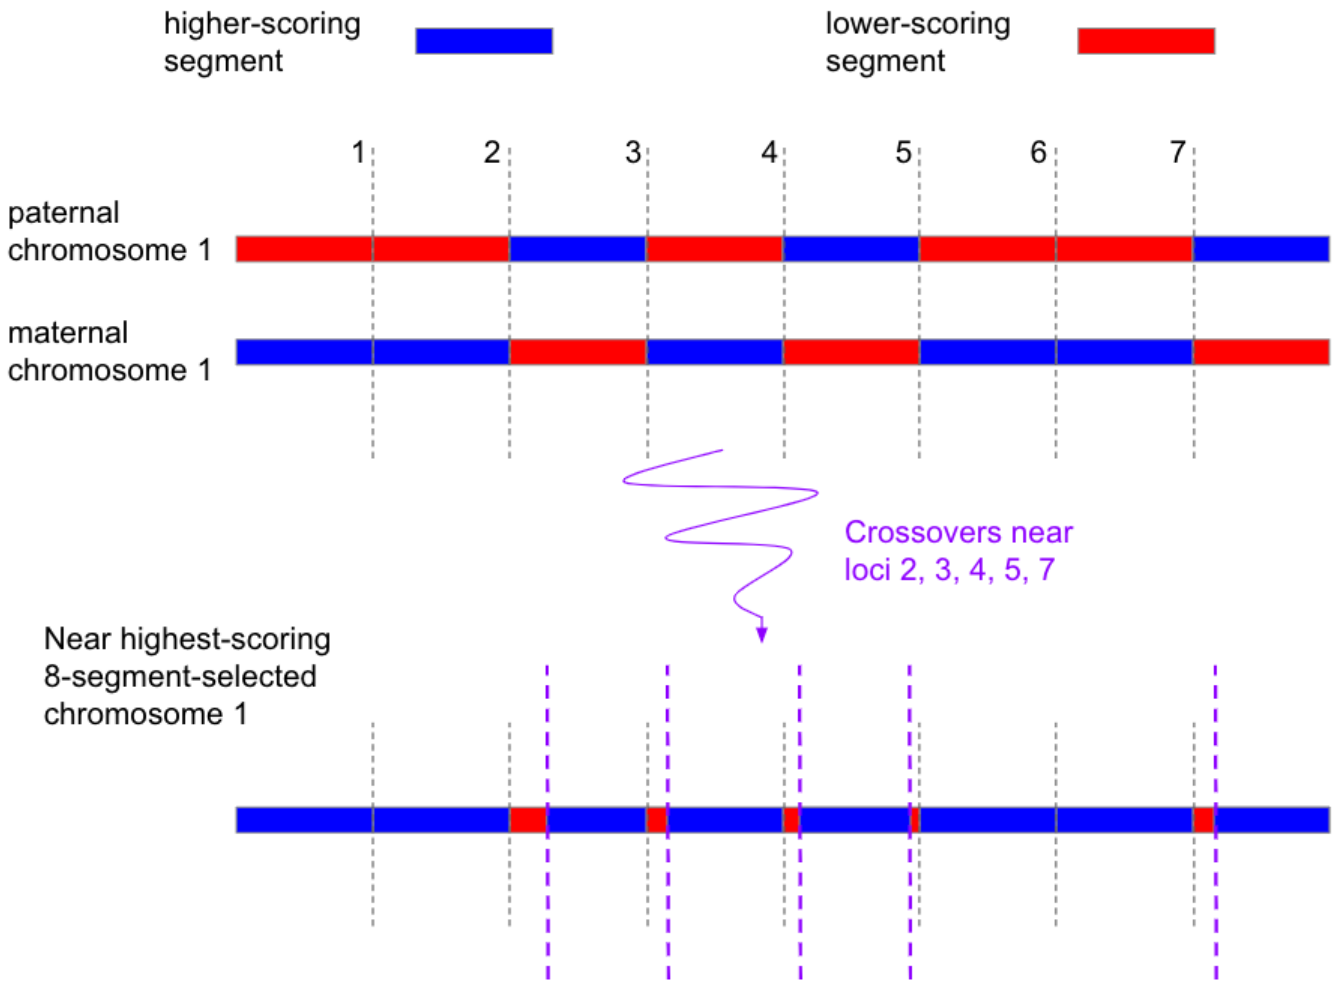
Some costs involved in IMS:

- Because you're trying to end up with a high-scoring cell, you want to maintain many positive variants in the population of cells. So you have lots of cell cultures at the same time, which takes money for equipment and culture media and labor.
- You have to do some form of DNA sensing to know something about the genomes of many cells, so you can select lines to propagate.
- You have to induce meiosis in very many cells. The process takes 2 or 3 weeks, and it will have some substantial attrition rate.
- You have to isolate and passage cells so you can proliferate lines as monoclones to sequence and select. (Well, not necessarily. You could do some more scalable process. E.g. you could use some fluorescent marker and FACS to find cells enriched in certain DNA segments.)
- All these processes can introduce mutations, e.g. cancer mutations, which you have to discard, adding more attrition.

Uncertainty in these costs makes it hard to estimate in advance the selection power per cost you can get with IMS. Aside from all these costs, the scheduling problem of when to combine, meiose, sequence, or discard which cells, is quite nontrivial; I don't understand it satisfactorily. It seems like it would be fun to build a proper simulator for the task, and then come up with and test strategies or do machine learning. Let me know if you're interested in doing that (it's not a priority for me).

The power of segmented selection

The simplest way to think about the power of IMS in general is to think of IMS as being "like chromosome selection, but for subdivisions of chromosomes". Instead of selecting the top-scoring chromosome 1, we select the top-scoring first half of chromosome 1 and the top-scoring second half of chromosome 1. We do this by arranging that the chromosomes with the appropriate crossovers (corresponding to our selection of segments between the crossover points) will make it into one cell together. An illustration:



If you double the number of segments that you select but keep everything else the same, how much does

selection power change? Each segment now has half the variance of the previous double-segments. So a segment SD is $1/\sqrt{2}$ as big as an old double-segment SD. But there are twice as many, and the benefits sum. So the overall effect is that the selection power goes up by a factor of $2/\sqrt{2} = \sqrt{2} \approx 1.4$.

In other words, in general to compute the power of segmented selection, we can just multiply simple selection from a gaussian by $\sqrt{k}$, where k is the number of (equal-length) segments.

For example, suppose we want to make a haploid cell from a single diploid. That means we'll get two choices for each segment; that means our base selection power is about 0.56 SDs.

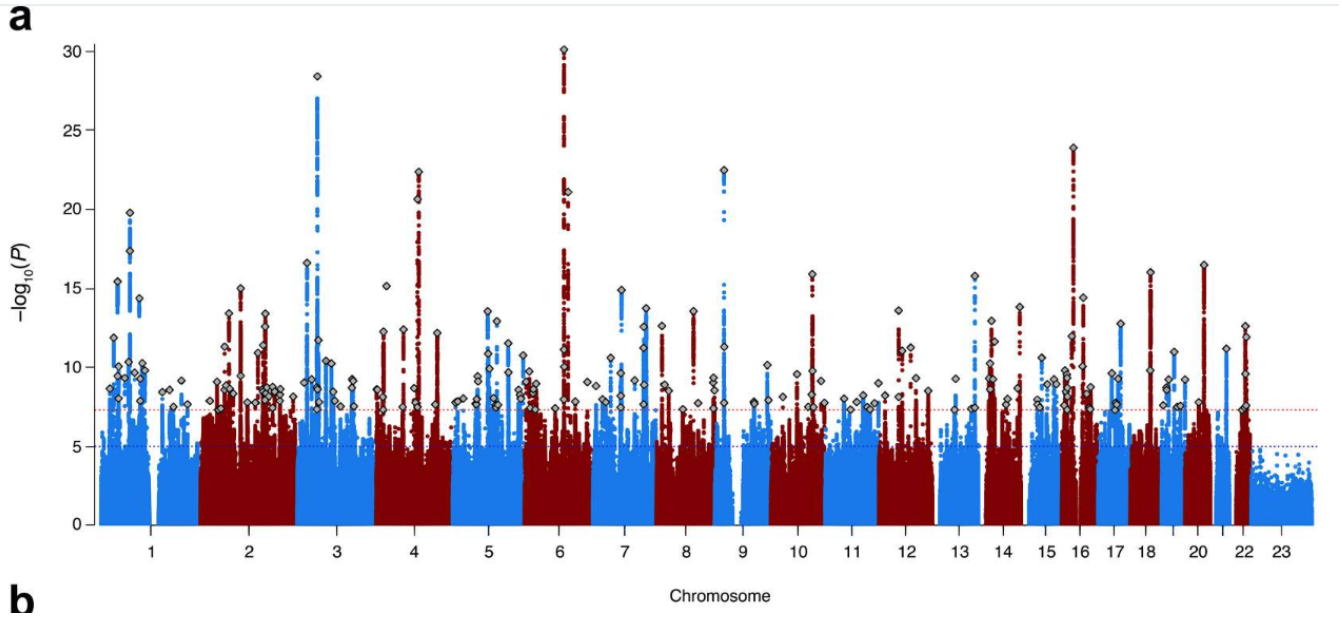
If we just do chromosome selection, that's 23 subdivisions. The total variance of the haploid is 1/2 the total population variance. So the scale factor is $\sqrt{1/2} \times \sqrt{23} \approx 3.4$. This says that chromosome selection should give $0.56 \times 3.4 \approx 1.9$ SDs, which is roughly what we got for haploid chromosome selection above. (The divergence is because here we used 0.56 for best of 2 from a standard gaussian, and because we're allowing selection on the 23rd chromosome.)

Now suppose we do segmented selection with 92 subdivisions. In other words, we're dividing each chromosome into 4 equal segments, and then picking the top-scoring of each segment. How strong is this? Multiplying the segments by 4 will multiply the power by $\sqrt{4}$, which I believe is 2. That gives $2 \times 1.9 = 3.8$ total raw SDs, or 4.5 after embryo selection, or 27 IQ points.

Breakdown at high granularity

How about 10 subdivisions? Now we get $1.9 \times \sqrt{10} \approx 6$ raw SDs. After embryo selection, that's about $6.7 \times .4 \times 15 \approx 40$ IQ points. However, at this level of granularity, these mathematical estimates may start to significantly break down in reference to actual genomes. The issue is that with 10 subdivisions of 46 chromosomes, we're looking at around 500 DNA segments in a whole genome. A PGS might only know about dozens or hundreds of SNPs. So it doesn't make sense to directly model a segment as having a score sampled from a gaussian, as many segments will only have a few or even no effectful SNPs. On the other hand, the variance is still there somewhere, so one could imagine having segments with uneven length but evenly distributed variance; I'm not sure exactly where things really break down.

That said, in the case of IQ, the trait is highly polygenic—maybe on the order of 10,000 variants are involved—and a substantial fraction of the trait variance is explained by a PGS, involving over a thousand variants. Furthermore, the SNPs associated with IQ do tend to be spread out all across the genome, rather than being highly concentrated in smaller regions. See for example:



(This figure shows p values and doesn't show effect sizes, but the point stands.)

(Figure 1a from Savage et al. 2018[170].)

Finally, in practice, parents will want to select for several traits. Effectively, they'll select for a combined PGS, which may be a weighted sum or a more complicated function of PGSes for diseases, health, personality,

[170]Savage, Jeanne E, Philip R Jansen, Sven Stringer, Kyoko Watanabe, Julien Bryois, Christiaan A de Leeuw, Mats Nagel, et al. 'Genome-Wide Association Meta-Analysis in 269,867 Individuals Identifies New Genetic and Functional Links to Intelligence'. Nature Genetics 50, no. 7 (July 2018): 912–19. https://doi.org/10.1038/s41588-018-0152-6.

IQ, or other traits. The combined PGS will have thousands of genetic variants to select for, spread across the whole genome. So these estimates of total raw SDs might mostly extend all the way to 10 segments per chromosome.

Segmented selection with one or more donors

If instead of making just one haploid, we make both, the gains add together, i.e. they are doubled: >50 IQ points for 4 subdivisions, or 80 IQ points for 10 subdivisions.

Say that instead of making a haploid from one parent, we make a haploid from two parents, to make a child with two genetic grandfathers and one genetic mother. That means that instead of sampling from 2 segments to get each 1 segment, we sample from 4 segments to get each 1. The expected maximum of 4 samples from a standard gaussian is about 1 SDs above the mean. Here's a snippet you can check these numbers with:

```
import numpy as np
def expected_max(test_runs=10_000_000, k=2):
    X = np.random.randn(test_runs, k)
    return np.mean(np.max(X, axis=1))
```

So with two parents to make a haploid, we get almost twice as much selection power as with one parent. For example, dividing each chromosome into 4 segments, we have 92 segments. The selection power is $\sqrt{1/2} \times \sqrt{1/92} \times 92 \times 1 \approx 6.8$ raw SDs. After embryo selection, that's about 55 IQ points. Doing this for both haploids is too many IQ points. Here's a table giving the numbers for segmented sampling of a standard gaussian, from here:

cells show SDs of maximum

number of subdivisions \ samples selected from	2	3	4	5	6	10	30	100	1000	10000	100000
2	0.8	1.2	1.46	1.64	1.79	2.18	2.89	3.55	4.58	5.45	6.2
3	0.98	1.47	1.78	2.01	2.19	2.67	3.54	4.34	5.61	6.67	7.59
4	1.13	1.69	2.06	2.33	2.53	3.08	4.09	5.02	6.48	7.7	8.77
5	1.26	1.89	2.3	2.6	2.83	3.44	4.57	5.61	7.25	8.61	9.8
6	1.38	2.07	2.52	2.85	3.1	3.77	5.0	6.14	7.94	9.43	10.74
10	1.78	2.68	3.26	3.68	4.01	4.87	6.46	7.93	10.25	12.18	13.86
23	2.71	4.06	4.94	5.58	6.08	7.38	9.8	12.03	15.55	18.47	21.03
46	3.83	5.74	6.98	7.89	8.59	10.44	13.85	17.01	21.98	26.12	29.74
92	5.41	8.12	9.87	11.15	12.15	14.76	19.59	24.05	31.09	36.94	42.05
138	6.63	9.94	12.09	13.66	14.89	18.08	24.0	29.46	38.08	45.25	51.5
184	7.65	11.48	13.96	15.78	17.19	20.87	27.71	34.01	43.97	52.25	59.47
460	12.1	18.15	22.08	24.94	27.18	33.0	43.81	53.78	69.52	82.61	94.03
1000	17.84	26.76	32.55	36.78	40.07	48.66	64.6	79.3	102.5	121.8	138.64
10000	56.42	84.63	102.94	116.3	126.72	153.88	204.28	250.76	324.14	385.16	438.43
100000	178.41	267.62	325.52	367.76	400.73	486.6	645.98	792.97	1025.03	1217.99	1386.44

One way to interpret these numbers is that the column labels (2,3,4,5,6,10,30…) say the number of copies of each chromosome you start with, and the values given in the cells show the SDs of power you get from making both haploids separately given that many chromosomes. So the 2 column is the most normal case; it shows two parents who each make a haploid from their own genome (which gives 2 chromosomes of each type) with segmented selection, and then combine them. The 4 column shows the case with two parents to make each haploid. (The odd-number columns are not very relevant in this context.)

The row labels (…,23,46,92,…) show the number of divisions of the whole genome that are involved (including the separation between chromosomes). So the 46 row shows diploid chromosome selection. The $184 = 4 \times 46$ row shows segmented selection with 4 subdivisions. As mentioned, a while past 184, things might start to break down.

Already with a normal, realistic case, IMS provides strong GV. In the 2 column (i.e. just one couple, two parents), and the 184 row showing 4 subdivisions, we get 7.6 raw SDs, or 46 IQ points.

The table can also be read as describing making one haploid; you just divide the values in the cells by $\sqrt{2}$.

The cost of poor man's chromosome selection

Keeping in mind that the optimal scheduling problem for IMS is complicated and hard to analyze, we can get still some sense for the problem. One angle of attack is to simplify way down like this: Assume that,

instead of meiosis, the division process simply picks one of two whole chromosomes at random.

This is like the problem of whole cell fusion and tetraploidy reduction, but with diploids and haploids rather than tetraploids and diploids. If you think about the problem a bit, I think you'll see that it's not so trivial.

It's easy to make some progress quickly: You just select haploids with a larger number of the higher-scoring chromosomes and combine them. This does work well, but it usually fails to get all 23 target chromosomes into one cell. The issue is that pretty soon, all the top-scoring cells are top-scoring because they came from the same earlier top-scoring cell, so they have a lot of overlap in the target chromosomes that they have, and leave some target chromosomes extinct. Of course, you can do more clever things, such as preserving rare top-scoring chromosomes, and trying to make pairs that stand to make the most gain (have the most disoverlap of target chromosomes). These work, but it's not straightforward to mathematically analyze.

There is, however, a "speed limit" we can describe. What's the best possible pairing of haploids to make? You always want to take your top-scorer, and then pair it with another equally high-scoring haploid, but with maximal disoverlap so that you have the most upside (in other words, the most variance in the number of target chromosomes in the child haploids). Suppose that magically you can always do this. How fast do you go?

There's another parameter to specify, which determines how good the next top-scoring haploid will be: how many haploids do you get from each cell? As a simplifying assumption, suppose you get enough samples so that you will gain $\sqrt{k}$ (i.e. one SD, ish) additional target chromosomes, above the mean, where k is the number of non-target chromosomes left in your top-scoring haploid.

So at each stage, you do $k_{t+1} := k_t - \sqrt{k_t}$. This converges quite quickly: $[11.5, 8, 5, 3, 1, 0]$. So, not all that much of a speed limit after all, but it can give us a target speed to aim for.

The cost of single-chromosome IMS

As a different approach to analyzing IMS, we can simplify the problem by assuming that we have a method for chromosome selection. Now the problem is: If the genome is just a single chromosome, how do you cheaply get a high-scoring recombinant chromosome? If you have a good answer to that, you apply it 23 times separately, in parallel, for each chromosome. Then at the end, you use chromosome selection to bring those chromosomes together.

Since whole cell fusion is a kind of chromosome selection, we've sort of broken the whole IMS problem into two independent pieces. You're weaving together chromosomes; and then you're putting them together. In reality, you wouldn't do these separately, because it's more efficient to do everything simultaneously "in superposition", i.e. both weaving high-scoring chromosomes and shuffling those chromosomes together into the same cell. So the costs don't directly add up.

How to do single-chromosome IMS? The simplest case is this:

- There is one chromosome type.
- Each chromosome is a pair (c_0, c_1), where the c_i are bits $\{0, 1\}$. The meaning of $(0, 1)$ is "this chromosome has a non-target first half and a target second half".
- Our goal is to get $(1, 1)$.
- We start with a diploid $((0, 1), (1, 0))$.
- We can sample haploids through "meiosis", where there's always exactly one crossover:
  - Randomly pick one of the two chromosomes.
  - Take some initial segment of the chosen chromosome, either nothing, the first bit, or both bits.
  - Fill in the remaining bits with the tail end of the non-chosen chromosome.

This case is easy to analyze. Half the time meiosis picks $(1, 0)$ to start with, and a third of the time it picks the middle for the crossover point. So if you sample about 6 haploids, you get $(1, 1)$.

What about if we divide the chromosome into more segments? How do you assemble the target segments into a single chromosome? E.g., suppose you start with

$$(0, 1, 0, 1, 0, 1, 0, 1)$$

$$(1, 0, 1, 0, 1, 0, 1, 0)$$

and you want to make

$$(1, 1, 1, 1, 1, 1, 1, 1)$$

How do you schedule pairings and meiosis to get the all-1s chromosome?

…

…

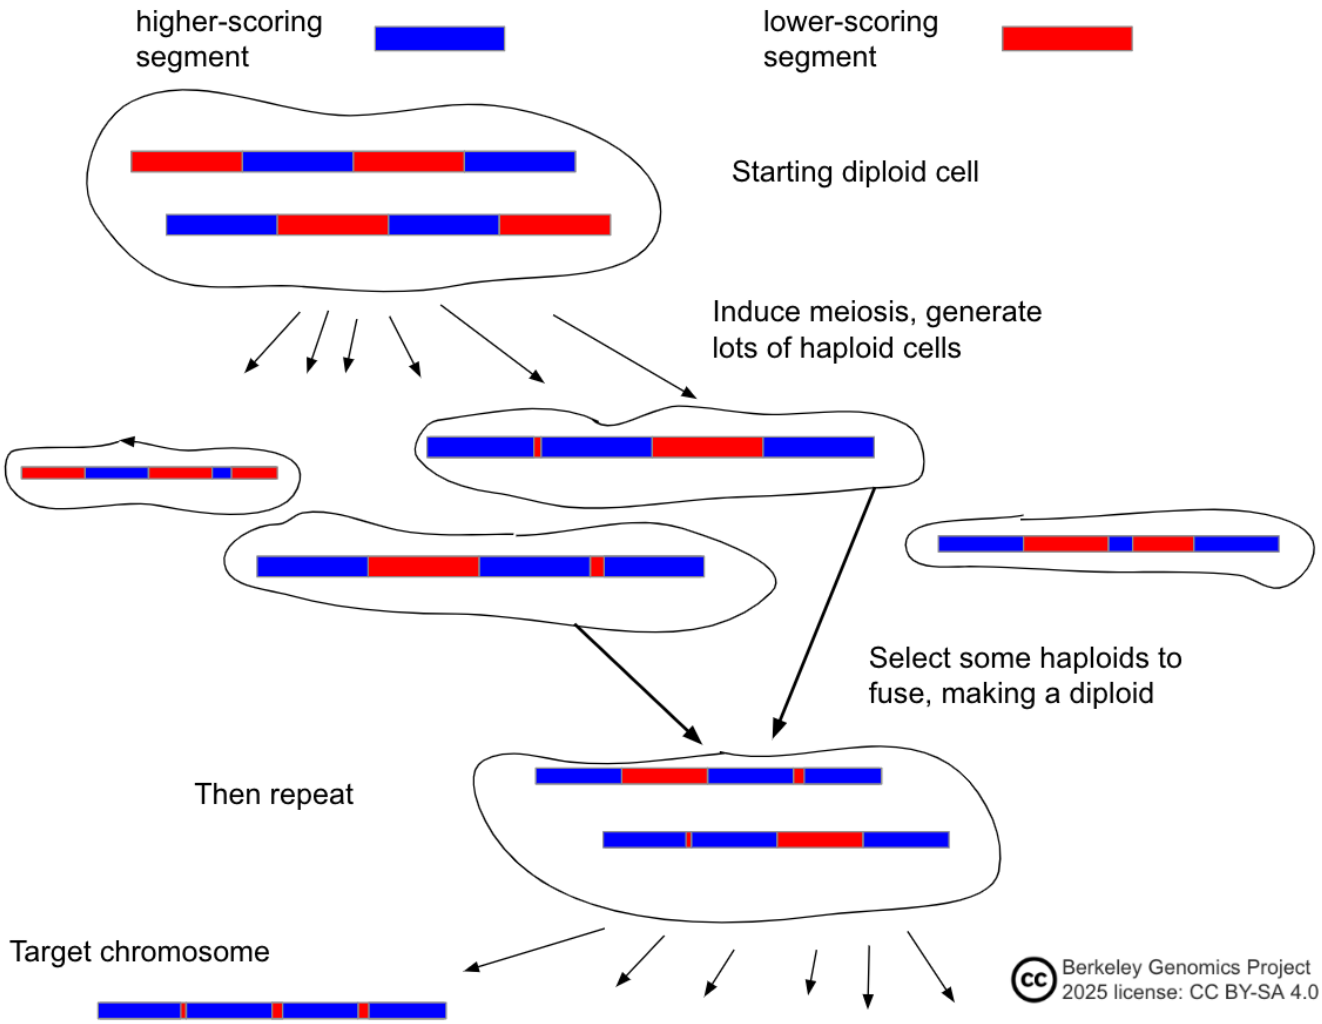
(Pause for thinking.)

…

…

We need a crossover at each comma, i.e. 7 total. By default, if you accept random crossover locations, the expected number of total crossovers increases by one each generation. So we'd need 7 generations. This isn't hard: In generation k you sample until you get a crossover at location k, adding 1s to the initial segment one by one in order. Surely we can do better?

One strategy you can follow is to assemble sub-chromosomes in a binary-tree pattern. We build up islands of 1s which double in size and halve in number each round. An illustration, in the smaller case of 4 segments rather than 8:



To step through in the case with 8 segments: In the first round, you sample enough haploids that you get every crossover point. Now you have

$$(1, 1, 0, 1, 0, 1, 0, 1)$$

$$(1, 0, 1, 1, 0, 1, 0, 1)$$

$$(1, 0, 1, 0, 1, 1, 0, 1)$$

$$(1, 0, 1, 0, 1, 0, 1, 1)$$

In the next round, we pair the first two of these and we pair the last two. We sample from each haploid until we get the desired crossovers, like this:

$$(1, 1, 1, 1, 0, 1, 0, 1)$$

$$(1, 0, 1, 0, 1, 1, 1, 1)$$

And then once again until we get the middle crossover.

This method takes roughly 8 samples for each pairing; it takes 3 serial generations (rounds of meiosis); requires proliferating 4 diploids; and requires producing and genotyping on the order of $18 \times (1+2+1) \approx 70$ haploids from diploids. If we did this 23 times and then did chromosome selection, we'd have to pay these costs 23 times (in parallel).

I would guess this strategy is basically optimal for the simple case we're dealing with.

However, in real life, there's a lot of utility being left on the table, potentially. All the assumptions we made are unrealistic, obscuring both obstacles and possible better strategies. E.g.:

- Sometimes meiosis produces more than one crossover in one chromosome. So potentially you could go faster.
- Chromosome crossovers don't just happen at your favorite 9 spots, they happen all over the place (though by no means uniformly distributed). This makes it more complicated to decide which haploids to keep; target segments could be broken up in whatever fraction. On the plus side, you might get especially helpful crossovers.
- We modeled the chromosome as being made of some equal-length segments with equal spread of scores between the two homologs. In real life, chromosome scores are more like a brownian motion (speaking imprecisely). They'll have lots of random little patches that are rather higher or lower scoring than usual. This means there's opportunity to get lots of especially good segments of various sizes, efficiently.
- We modeled a single chromosome. In real life there are 23 chromosomes. We could, to a significant extent, parallelize the individual chromosome-segment-selection processes within one cell. In other words, by causing a single diploid cell to undergo meiosis, we get 23 crossovers, one in each chromosome; it may be that several of them are useful.
- We assumed that we start with a single diploid cell, and want to produce a single haploid cell. In real life, we might start with cells from several different donors, and then produce a diploid or haploid. This is more complex, but the additional genetic variants raise the achievable ceiling.
- We assumed we genotyped all the haploids. Ideally we'd use a much cheaper DNA sensing method. But either way, it's possible to not genotype every haploid, and/or not wait for genotypes to be processed before proceeding to the next steps. This is sort of like branch prediction in a CPU. E.g. we could start inducing meiosis in a diploid before we've gotten back the sequencing results, and then later scrap the resulting haploids if we hear back about the diploid and decide it won't produce the haploids we want. Also, we could use multiple sensing methods, e.g. very coarse FACS to pick generally target-enriched candidate haploids to make monoclones from, and then do some cheap SNP array to learn more about those monoclones before pairing.
- In real life, we have to at least sometimes do whole genome sequencing to screen for de novo mutations.
- In real life there are many decisions to be made about which cell lines from earlier stages of the process to continue culturing, in order to preserve variation, and when to bring them back to reintroduce variants.
- We considered 8 segments, a power of 2. But 6 or 10 segments might be more favorable. The simple binary tree doesn't work in this case.

Appendix: Detailed estimation of embryo selection

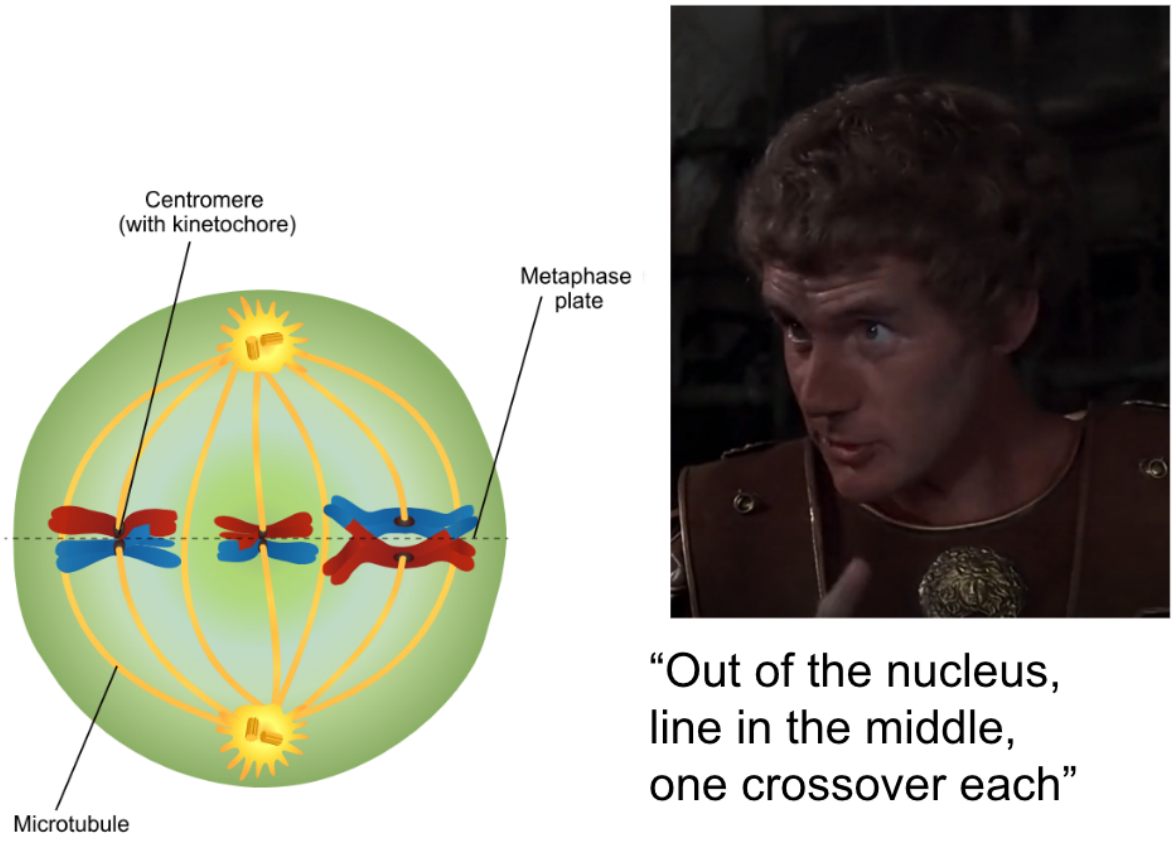
Note: For most purposes, this section gives too much detail. To just understand the expected gains from simple embryo selection, it suffices to look at some gaussian and ask about the order statistics. But, I wanted to understand the error bars a little bit.

Here's the code for simulating random genomes (modeled as subdivided gaussians): https://gist.github.com/tsvibt/89cec1b4fd7d54be04ba34a0059bcc4c

This isn't all of the code used to make the graphs below, but it's the core model used, so if you want to check its correctness or play with it, that's the place to start.

Simple embryo selection simulations

In what follows: Chromosomes are modeled as being made of 100 segments, with each segment sampled from a gaussian with 1/100 of the chromosome's variance, which in turn is proportional to the chromosome's share of total genome length. Meiosis uniformly at random picks a point along the chromosome, and a parity (which chromosome first), and puts the result as the child chromosome; one crossover each.



(Modified from https://en.wikipedia.org/wiki/File:Meiosis_Stages.svg.)

Analysis of simple embryo selection

We want to understand the effect of choosing the highest-scoring embryo that one couple produces. We can break this up into two questions:

- What is the distribution of embryos from one couple? In other words, what is the distribution of the score of embryos, sampled by combining two parent genomes randomly in the way that natural reproduction combines them?
- Given the distribution of embryos from one couple, what is the distribution of the highest scoring embryo out of a sample of k embryos from the couple?

What is the distribution of embryos from one couple?

For each parent, for each chromosome, we sample a chromosome, independently of the other chromosomes. (That is an assumption about the biology—that crossover points and assortment aren't correlated across chromosomes.) Sampling a chromosome of index i means sampling a recombination between the parent's

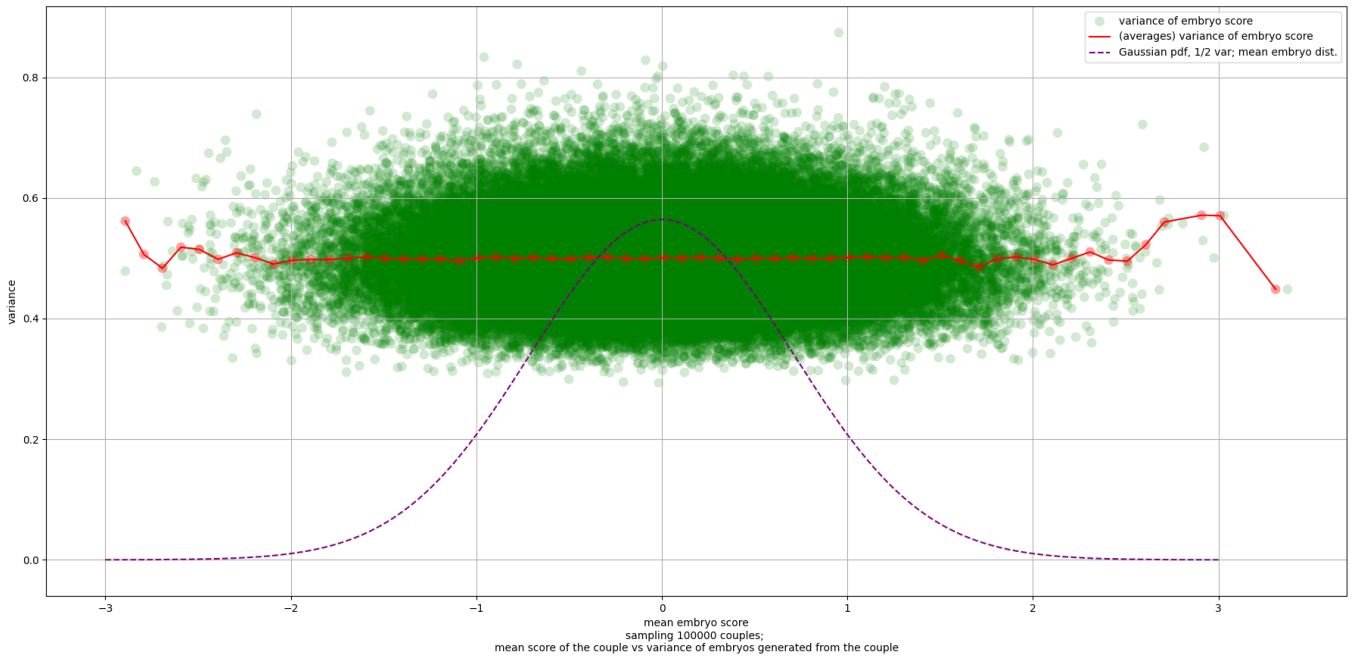
two chromosomes of index i. Modeling the distribution for one chromosome isn't trivial; see the appendix "Best crossover". But we can zoom out and just model the haploid as a sum of similarly-distributed variables. This gives approximately a gaussian, and then the embryo is just a sum of the two haploids, so it's another gaussian.

Now we just need to know the mean and variance of the gaussian. The mean is just the mean score of the two parents's genomes. What is the variance?

Heuristically speaking: The variance of all people is the same as the variance of all embryos. Embryos are sampled by first sampling a couple, and then sampling an embryo from that couple. The mean of the embryos from a couple is the mean of the couple's genomes. The score of an embryo from a couple is the mean of the couple, plus a sample from the couple's embryo-variance. So the overall variance of all embryos is the variance of the mean of a couple, plus the couple's embryo-variance. (This step needs more precision. It uses linearity of variance for the sum of independent variables. It also uses that conditioning on the couple's mean being some k_1 produces a distribution of embryos that's the same as conditioning on the mean being k_2, but with means shifted up or down. The conditional distribution introduces small correlations between all the values of segments of chromosomes in both parents, but the point is that the shape of the conditional distribution (ignoring means) doesn't depend on the value being conditioned on; see this stackexchange answer.)

The couple's mean score is half the sum of their individual scores. Variance is additive, so $\mathrm{Var}(G_1 + G_2) = 2\mathrm{Var}(G_1)$. But halving the score will 1/4 the variance (the score is in SD-space, the variance is in SD^2-space). So the variance of the couple's mean score is half the variance of the population. That leaves the other half of the overall variance of all embryos to be generated by a single couple's embryos. So a couple's embryo-variance is around half the population variance. (Again, this is not precise.) In other words, the typical couple's embryos are scores distributed as a gaussian with standard deviation roughly equal to $1/\sqrt{2}$ the standard deviation of the population.

We can check this with simulations. Here's a graph plotting the variance of embryos from couples. (That is, the variance of total score, in the SD-scale of the population.) The green dots are couples. The x position of a point shows the mean of the couple's genomes, i.e. the mean of the couple's embryos. The y position of a point shows the variance of embryos sampled from the couple. The red line shows the averages of variances (binned together in some range on the x-axis). The gaussian is the probability density function (pdf) of a couple having a given mean score (with probability density on the y axis).



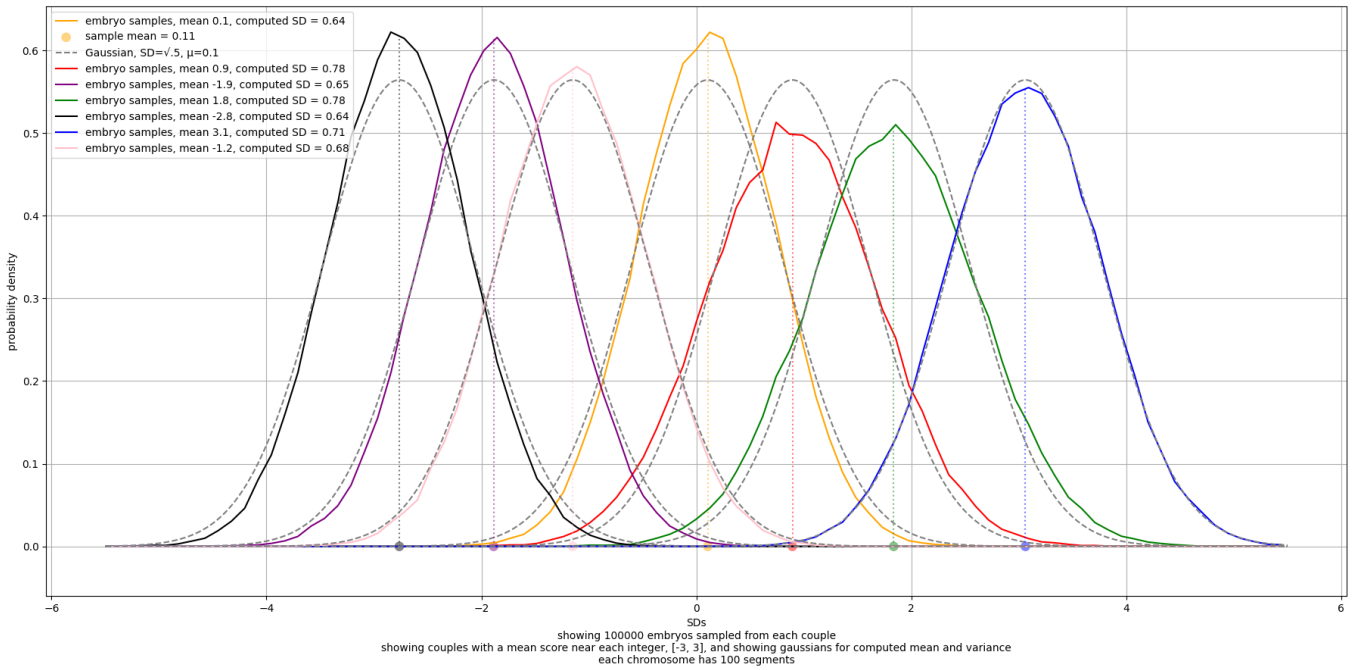
We see a few things:

- The variance is about 1/2.
- The variance does not change for couples far from the mean genome.
  - (In theory, that is. In theory in theory, you never use up variance. In theory in practice, selecting very hard would use up variance—the genome isn't actually an infinitely divisible Weiner process, it's a finite list of base pairs. But in practice in practice, you don't use up much variance. E.g.

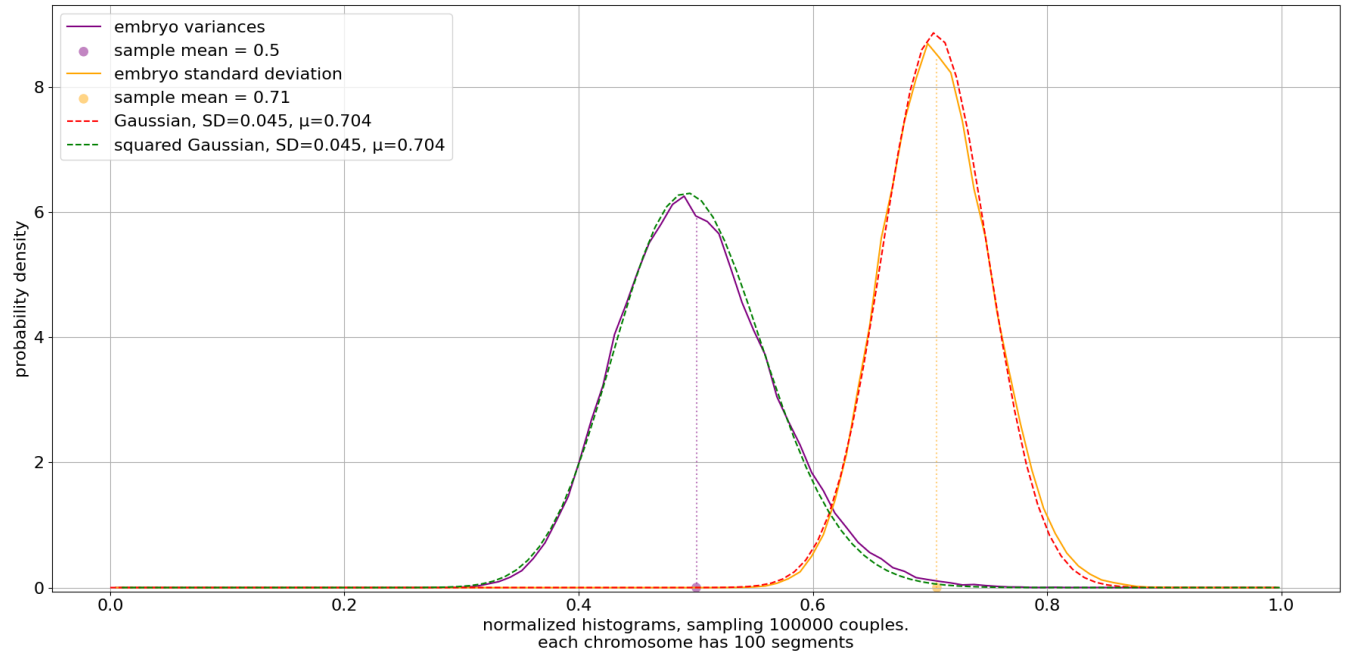
if intelligence has on the order of 10^4 relevant alleles, one SD is 50 alleles. So you could select by 20 SDs, i.e. 1000 alleles, which is way too far anyway, and still have only used up less than half of the available variance. See the limits section of "The Power of Selection".)

- There is some variance in the variance of a couple's embryos. It looks like the distribution is a little skewed to be higher. It looks like the range is roughly ± 0.1. Taking square roots to translate to SD-space, we have roughly 0.7 ± 0.07. So selection will work similarly well for everyone, but could have effects that vary by about 10%, compared to a modal couple with embryo variance .5.

For the sake of concreteness, below is a plot that shows seven couples. Each couple then generates many embryos, and we graph a histogram of the embryos's scores. For each couple we also graph a gray dotted gaussian with $SD = 1/\sqrt{2}$, centered around the couple's computed mean. Comparing the gaussian gives some sense of the variation in the variation; the couples are all quite close to their gaussians, but there's some noticeable differences.



To get efficient estimates for sampling maximum embryos from couples, we want a better estimate of the parameters of the distribution of embryo-distributions. Below is a graph showing a histogram of, for each of 30,000 couples, the variance and standard deviation of embryos generated by that couple (computed by aggregating the couple's haploid variances). We see that the SD of the SD is about .045. The distribution is not a gaussian (e.g. because it has probability 1 of being nonnegative), but empirically it's pretty close to one.



The "squared gaussian" is the distribution of the square of the SD value. The formula used is

$$\frac{\varphi\left(\sqrt{x} \mid \sigma^2, \mu\right) + \varphi\left(-\sqrt{x} \mid \sigma^2, \mu\right)}{2\sqrt{x}}$$

where

$$\sigma \approx .045, \mu \approx \sqrt{.496}$$

by eyeballing it. Here $\varphi\left(z \mid \sigma^2, \mu\right)$ is a gaussian distribution. I think this is a noncentral chi-squared distribution.

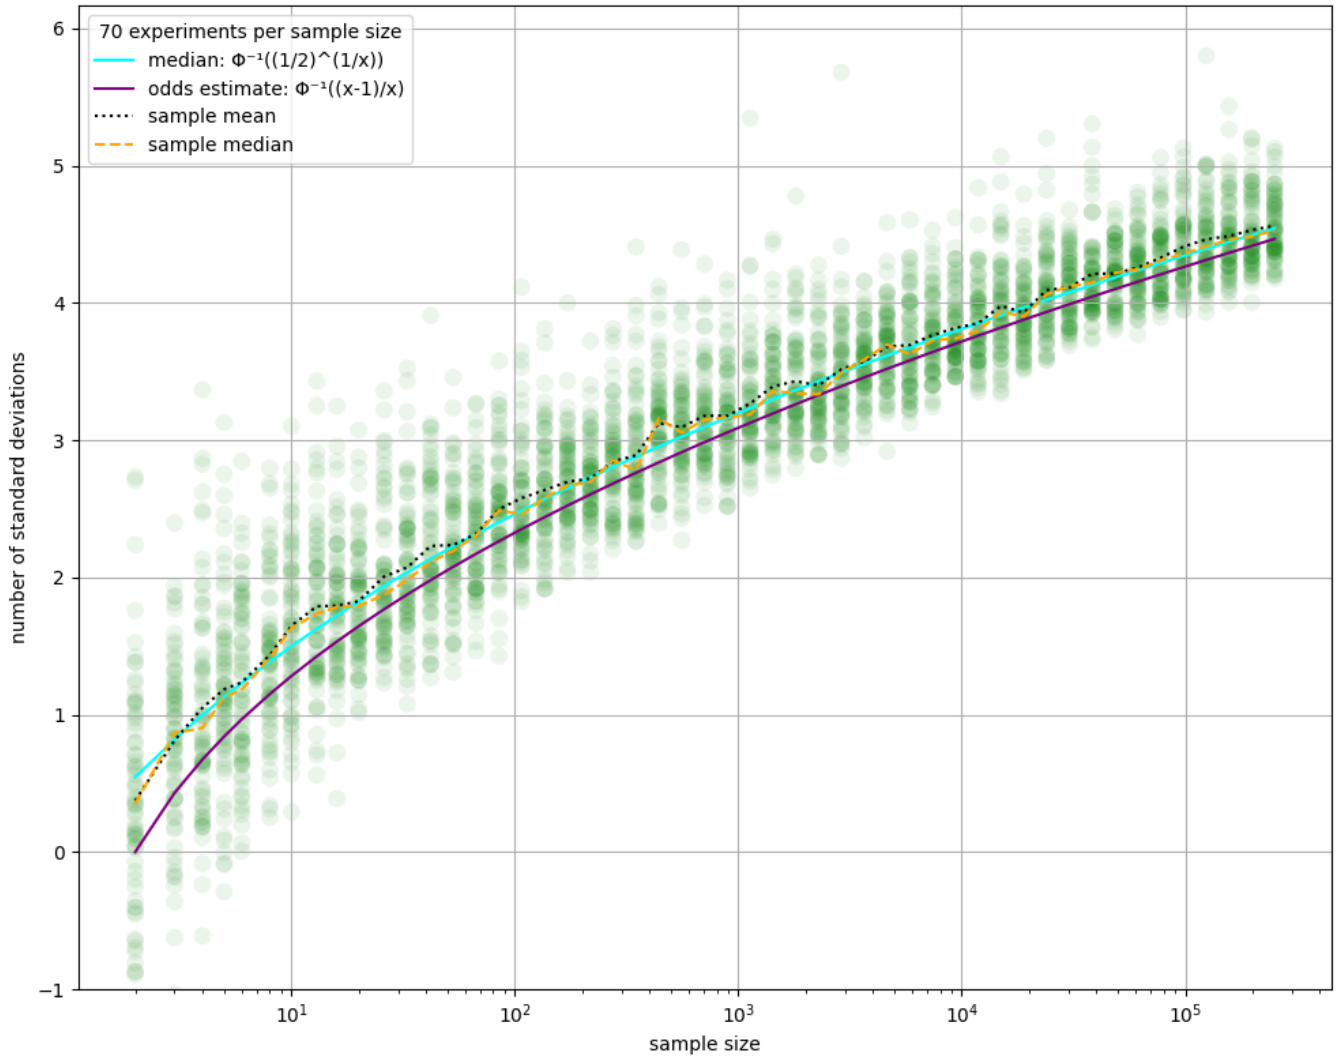
What is the maximum scoring embryo from one couple?

We have some handle on the distribution, across couples, of the distribution of a couple's embryos. The embryo-distribution is roughly a gaussian. (We can basically ignore the mean because, as we saw, the embryo distribution doesn't appreciably change with the couple's genome mean. This relies on PGSes having many hundreds of target alleles, which is already true as of 2024.) The embryo-distribution gaussian has SD sampled from roughly a gaussian with mean $\approx .704$, SD $\approx .045$.

So, to model a couple, we first sample from the embryo-distribution-SD gaussian, to get an SD σ. Then, to sample embryos, we sample from a gaussian with mean 0 and SD σ. This represents not the raw score of the embryo, but rather the difference between the embryo and the couple's mean score. To further standardize, we can instead model sampling from a standard normal distribution (mean = 0, SD = 1). This represents "couple SDs", i.e. SDs of this specific couple's embryos. Then, whatever the maximum is, we can convert that back into real SDs (i.e., SDs of the general population) by multiplying by σ, and then adding the couple's mean.

Now we just need to know the maximum sample from a standard normal. For this, see the Simple Selection section of "The power of selection".

The basic answers we want can be read off the last two graphs in the "Maximum sample out of n" section. Those graphs show the value of the maximum sample from a standard normal, with the number of samples maxxed over shown on the x-axis and the SD value shown on the y-axis. Here's the logscale one:



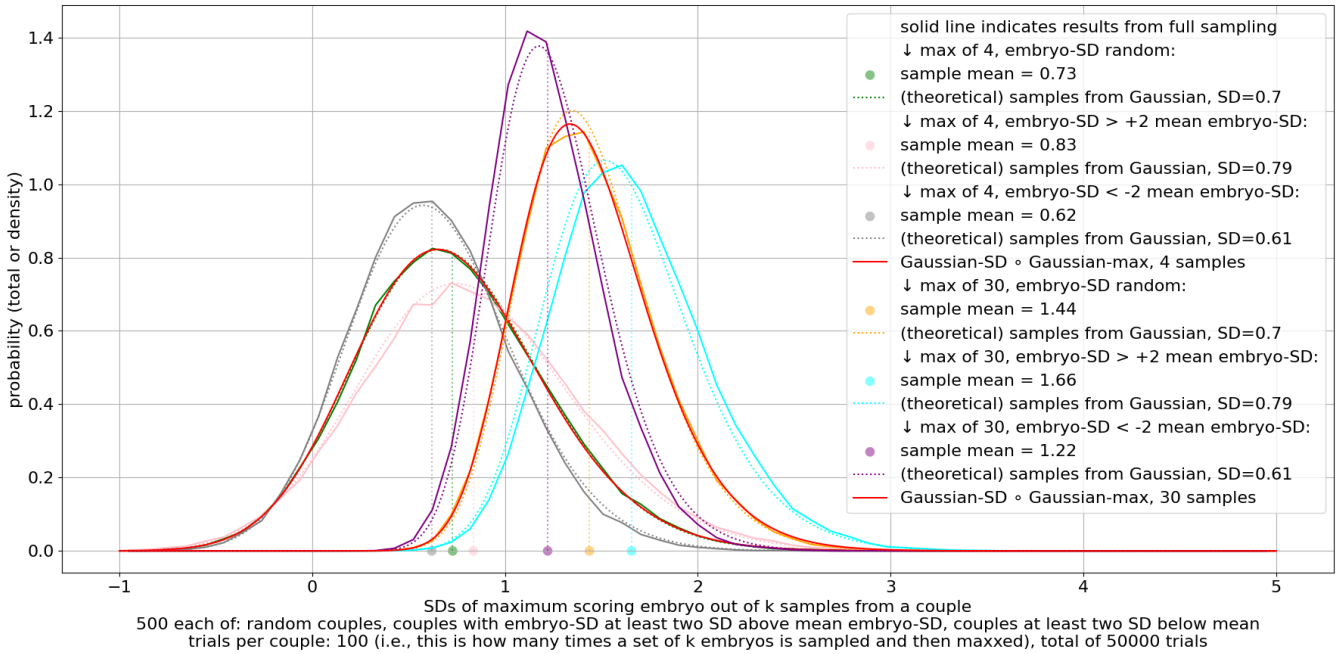
To convert to the values for one couple's embryos, we have to adjust to the couple's embryo-distribution. In the modal case, the couple's embryos are distributed like a gaussian with SD about .7. So we have to multiply the values we read on those graphs by .7. E.g. sampling about 30 embryos would give around $2 \times .7 = 1.4$ SDs of selection power. This says how much higher the embryo scores, compared to the mean of the parents's scores.

Looking at the logscale table, and considering a very optimistic case where IVO is available, we see that sampling 100,000 embryos would give roughly $4.5 \times .7 \approx 3$ SDs of selection power. However, this is only raw selection power. The resulting embryo is 3 SDs extreme, as measured by whatever we were selecting for. But our PGS is only partially correlated with the trait/s we're selecting for. As a toy example, suppose we only select for IQ, and we have a PGS that correlates .5 with the trait (currently there are IQ PGSes that correlate around .3 or .4). So we multiply the raw selection power of 3 SDs by .5 to get 1.5 SDs of effective selection power. Thus, even with very many embryos (requiring IVO), simple embryo selection only gets twenty-something IQ points. (And in reality you would use some selection power to select for other traits, e.g. to decrease disease risk.)

These estimates are reasonable rule-of-thumb estimates. But it would also be good to understand a bit more about the distribution of maximum scores. Reason 1: To get a sense of how much variation to expect from a real-life implementation, e.g. to not be too disappointed with an outcome a little below the mean outcome. Reason 2: Some of this understanding will transfer to understanding the results of other sampling procedures, e.g. sampling haploids.

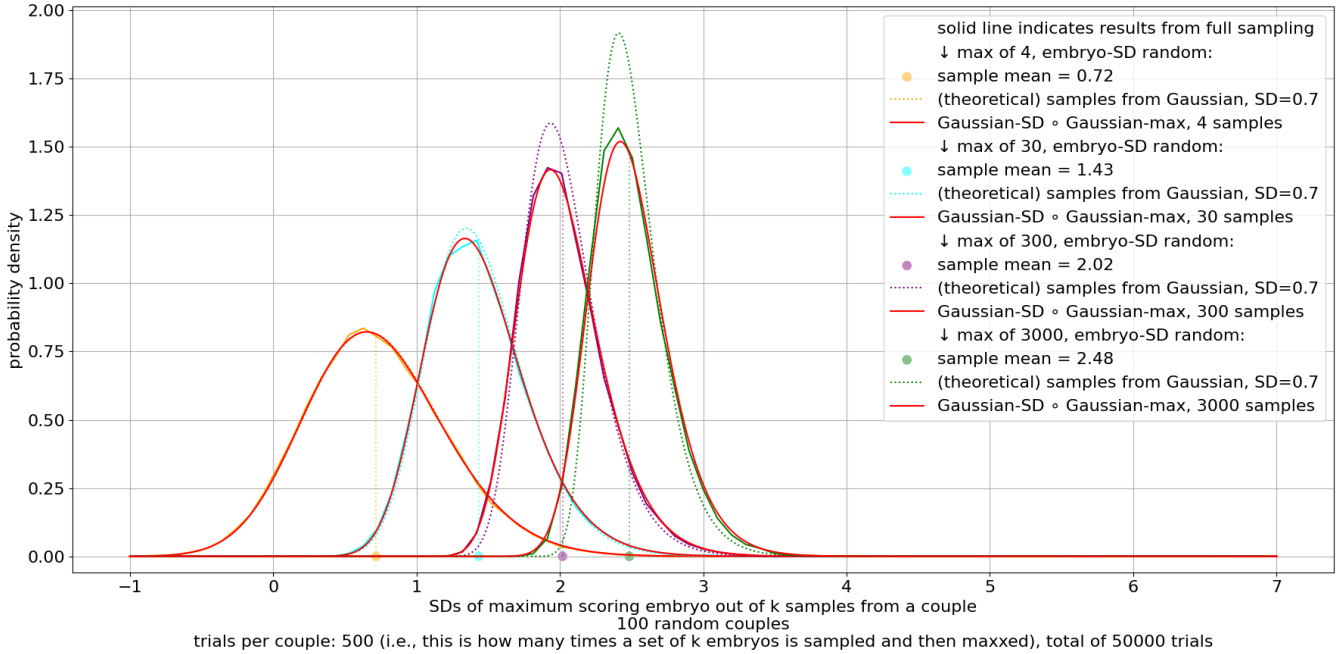
Here's a graph showing the results of taking a couple, getting some embryos from them, and then taking the highest-scoring one. The graph shows the results from taking 4 embryos, or taking 30 embryos; and shows the results for a set of random couples; and for a set of couples selected to have embryo scores distributed scores with standard deviation more than two embryo-SD-SDs greater than the mean embryo-SD ($.7 + 2 \times .045 = .79$); and for a set with embryo-SD less than two embryo-SD-SDs under the mean embryo-SD ($.7 - 2 \times .045 = .61$). (We're always looking at the difference between the embryo score and the couple's mean score; or in other words, we're conditioning on all the couples having mean score 0 SDs. These distributions are basically the same for any couple, since embryo variance isn't much affected by the

couple's scores for highly polygenic traits.)



We see that there's substantial variance in the max. This persists even for very large sample counts.

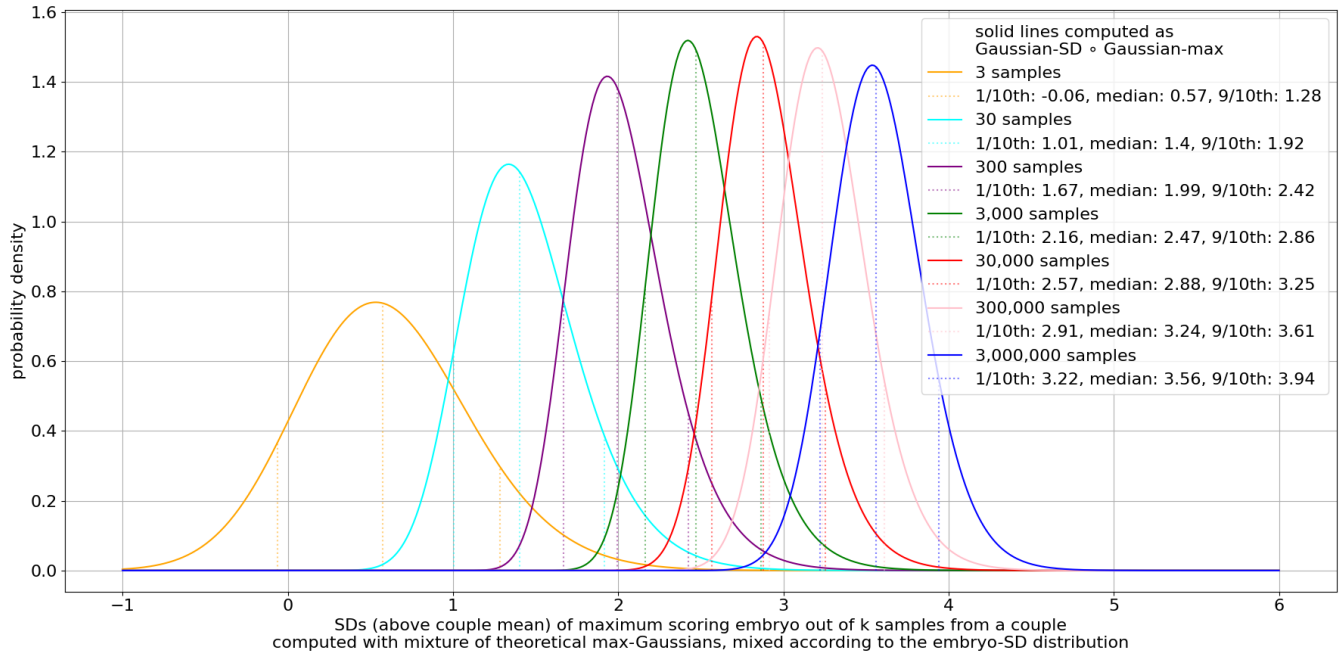
It's computationally expensive to do this sampling for large sampling counts. But we can use the theoretical predictions, described in the "Maximum sample out of n" section. The simplest way is to use the max-gaussian distribution, for a gaussian with SD equal to the mean embryo-SD, .7. This works reasonably well; these are the dotted lines in the above graph. For larger sample counts, this starts to diverge a little from the sampled results. We can see this by using another cheap approximation: Instead of actually sampling embryos, we can use the approximation of the embryo distribution as a gaussian. So we sample couples and compute their embryo SD, and then to simulate sampling embryos, we sample from a gaussian with that SD and with mean 0. Here's a graph showing this sampling, and showing the simple max-gaussian distributions:



We see that for larger sample counts, the simple max-gaussian is significantly tighter of a distribution compared to the actual samples. That's because the actual samples are the max from gaussians, but the gaussians have SD that has variance. We can better approximate the distribution by modeling that process. Instead of using one max-gaussian with mean SD, we take a mixture of max-gaussians. We weight the max-gaussian for σ according to the probability of σ in the gaussian that describes embryo-SDs, i.e. mean $\approx .704$, SD $\approx .045$. The red lines in the above graphs show a rough approximation of this distribution (a mixture of 100 max-gaussians). We see that it's a very close fit, practically, and it's much faster to compute

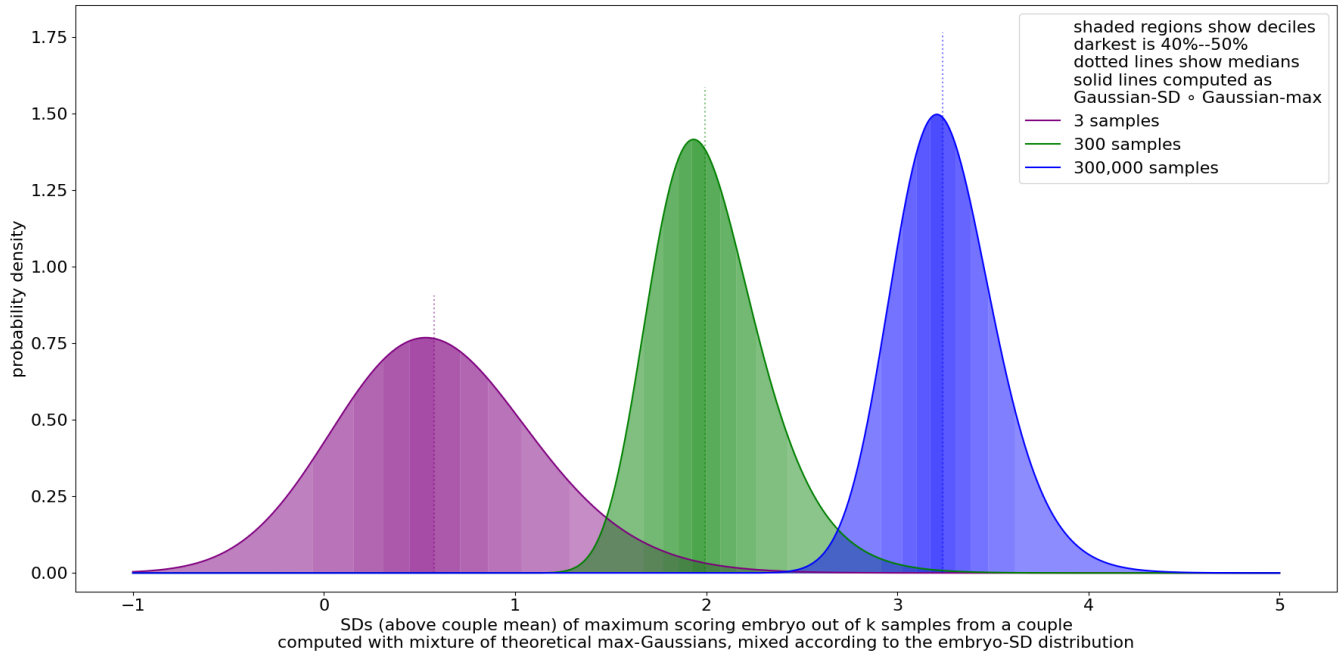
than sampling.

Using this mixture, we can see the results of maxxing over a set of embryos from one couple, for very large numbers of embryos:



This plot shows the lower decile, the median, and the upper decile, as dotted lines; and shows the mixed-max-gaussian distribution for various 3×10^k. (Notice that the upper decile, the rightmost of the three dotted lines in each color, for the max-3×10^k distribution tends to be close to the median line (the middle of the dotted lines) for the max-$3 \times 10^{k+1}$ distribution. This isn't a coincidence: The upper decile is about 1 in 10 rare, so getting an upper decile outcome happens roughly when you sample ten times. Sampling 10 times from the max-3×10^k distribution is equivalent to sampling once from the max-$3 \times 10^{k+1}$ distribution.)

Plotting with deciles shaded, to get a sense for the dispersion:



Putting together our numbers in a table:

first row shows distribution used. second row shows metric (decile or expectation).
$\mathcal{N}(\mu,\sigma)$ is a normal distribution with mean μ and SD σ. [ϕ]k means the distribution of X(k), the kth order statistic, i.e. the largest value of k samples from ϕ
$\mathbb{E}\_\sigma$ [$\mathcal{N}(0,\sigma)$]k means the distribution from taking the mixture of [max of k samples from $\mathcal{N}(0,\sigma)$], σ sampled from embryo SD distribution $\mathcal{N}(.704,.045)$

number (k) of embryos sampled. eX means 10^X

	[$\mathcal{N}(0,1)$]k	[$\mathcal{N}(0,.647)$]k	[$\mathcal{N}(0,.762)$]k	[$\mathcal{N}(0,.704)$]k				$\mathbb{E}\_\sigma$ [$\mathcal{N}(0,\sigma)$]k			
	expectation	expectation	expectation	expectation	1/10	median	9/10	1/10	median	9/10	median IQ points
1	0.00	0.00	0.00	0.00	-0.90	0.00	0.90	-0.90	0.00	0.90	0
2	0.56	0.36	0.43	0.40	-0.34	0.38	1.15	-0.34	0.38	1.15	2
3	0.85	0.55	0.64	0.60	-0.06	0.58	1.28	-0.06	0.57	1.28	3
4	1.03	0.67	0.78	0.72	0.11	0.70	1.37	0.11	0.70	1.37	4
5	1.16	0.75	0.89	0.82	0.24	0.80	1.43	0.23	0.79	1.44	5
6	1.27	0.82	0.97	0.89	0.33	0.87	1.49	0.33	0.86	1.49	5
7	1.35	0.87	1.03	0.95	0.41	0.93	1.53	0.41	0.92	1.54	6
8	1.42	0.92	1.08	1.00	0.47	0.98	1.57	0.47	0.97	1.57	6
9	1.49	0.96	1.13	1.05	0.53	1.02	1.60	0.53	1.01	1.61	6
10	1.54	0.99	1.17	1.08	0.58	1.06	1.63	0.57	1.05	1.64	6
e1.5	2.06	1.34	1.57	1.45	1.04	1.42	1.91	1.02	1.42	1.93	9
e2.0	2.51	1.62	1.91	1.77	1.41	1.73	2.17	1.38	1.73	2.19	10
e2.5	2.89	1.87	2.21	2.04	1.72	2.01	2.40	1.68	2.01	2.44	12
e3.0	3.24	2.10	2.47	2.28	2.00	2.25	2.61	1.94	2.25	2.66	14
e3.5	3.56	2.30	2.71	2.51	2.24	2.48	2.81	2.17	2.48	2.87	15
e4.0	3.85	2.49	2.93	2.71	2.47	2.68	3.00	2.38	2.69	3.07	16
e4.5	4.13	2.67	3.14	2.91	2.68	2.88	3.17	2.57	2.89	3.26	17
e5.0	4.38	2.83	3.34	3.09	2.87	3.06	3.34	2.75	3.07	3.44	18
e5.5	4.63	2.99	3.53	3.26	3.05	3.23	3.50	2.92	3.24	3.62	19
e6.0	4.86	3.14	3.71	3.42	3.23	3.40	3.65	3.08	3.41	3.79	20
e6.5	5.09	3.29	3.88	3.58	3.39	3.56	3.80	3.23	3.57	3.95	21
e7.0	5.30	3.43	4.04	3.73	3.55	3.71	3.95	3.38	3.72	4.11	22

cells show SDs (or IQ points)
.647 and .762 are 1/10th and 9/10th decile values for the embryo SD; .704 is the mean embryo SD
IQ points assumes a PGS with R of .4
IQ = (SDs of raw selection power) × .4 (effective SDs / raw SD) × 15 (IQ points / effective SD)

This shows the SDs above the couple mean of the highest-scoring embryo of k embryos sampled from a couple, with various approximations. The first few columns just show normal distributions $\mathcal{N}(0,\sigma)$ with SD σ being 1, or the mean embryo SD .704, or the bottom or top decile embryo SDs. These distributions represent the distribution for a couple with the corresponding embryo SD, after sequencing the parents and calculating that embryo SD. (The various $\mathcal{N}(0,\sigma)$ are just scaled versions of $\mathcal{N}(0,1)$, so their statistics are also just scaled versions. E.g. a number in the expectation column for $\mathcal{N}(0,.704)$ is just .704 times the corresponding number in the $\mathcal{N}(0,1)$ column. This is not true for the mixture distribution.)

The distribution labeled $\mathbb{E}\_\sigma[\mathcal{N}(0,\sigma)]k$ is

$$\mathbb{E}_{\sigma\sim\mathcal{N}(.704,.045)}[\mathcal{N}(0,\sigma)]^k$$

i.e. the mixture of k-max-embryo distributions, where σ is sampled from the embryo-SD distribution $\mathcal{N}(.704,.045)$. This represents the distribution for a couple's ultimate outcome from maxxing over k embryos, before we know the parents's genomes—i.e., we sample a couple from the population and then also sample k embryos and take the highest scoring.

Takeaways:

- For most purposes, we can just use the mixture distribution $\mathbb{E}_\sigma[\mathcal{N}(0,\sigma)]^k$. The median of the mixture distribution is similar to the mean and to the central tendency of the narrower $[\mathcal{N}(0,.704)]^k$ distribution that uses the mean embryo-SD. The mixture is wider.
- For convenience we can just use the median as our estimate. It's easier to compute, it's a slightly conservative estimate, and it's more comparable to the deciles. But the k-max distribution is skewed high, so the expectation is a little higher than the median.
- We quickly hit a big wall when sampling more and more. Going from 1 to 100 embryos gives +10 IQ points; going from 100 to 1,000,000 also gives +10 IQ points.
- The IQ estimates, or similar estimates for any other trait, depend on the PGS for that trait. If the PGS has an R value of .6 instead of .4, you get +50% the effect.

To get a better sense of the range of outcomes, we can zoom in on this mixture distribution:

all estimates from $\mathbb{E}_\sigma[\mathcal{N}(0,\sigma)]$k, which is the distribution from taking the mixture of [max of k samples from $\mathcal{N}(0,\sigma)$], σ sampled from embryo SD distribution $\mathcal{N}(.704,.045)$
M = median, 80% = 9/10 decile - 1/10 decile i.e. the width of the central 80% range

number (k) of embryos sampled. eX means 10^X

	1/10	M	9/10	M - [-.5+1.6$\sqrt{(\log10(k))}$]	M - (1+log10(k))/2	80%	80% - [1-.25log10(log10(k))]	(1/10)-M	(9/10)-M	(1/10-M)/80%	(9/10-M)/80%
1	-0.90	0.00	0.90	0.50	-0.50	1.80	-inf	-0.90	0.90	-0.50	0.50
2	-0.34	0.38	1.15	0.00	-0.27	1.49	0.36	-0.72	0.77	-0.48	0.52
3	-0.06	0.57	1.28	-0.03	-0.16	1.35	0.27	-0.64	0.71	-0.47	0.53
4	0.11	0.70	1.37	-0.04	-0.10	1.26	0.21	-0.59	0.67	-0.47	0.53
5	0.23	0.79	1.44	-0.05	-0.06	1.21	0.17	-0.56	0.65	-0.46	0.54
6	0.33	0.86	1.49	-0.05	-0.03	1.16	0.14	-0.53	0.63	-0.46	0.54
7	0.41	0.92	1.54	-0.05	0.00	1.13	0.11	-0.52	0.61	-0.46	0.54
8	0.47	0.97	1.57	-0.05	0.02	1.10	0.09	-0.50	0.60	-0.45	0.55
9	0.53	1.01	1.61	-0.05	0.04	1.08	0.08	-0.49	0.59	-0.45	0.55
10	0.57	1.05	1.64	-0.05	0.05	1.06	0.06	-0.48	0.58	-0.45	0.55
e1.5	1.02	1.42	1.93	-0.04	0.17	0.90	-0.05	-0.39	0.51	-0.44	0.56
e2.0	1.38	1.73	2.19	-0.03	0.23	0.81	-0.12	-0.35	0.46	-0.43	0.57
e2.5	1.68	2.01	2.44	-0.02	0.26	0.75	-0.15	-0.32	0.43	-0.43	0.57
e3.0	1.94	2.25	2.66	-0.02	0.25	0.72	-0.16	-0.31	0.41	-0.43	0.57
e3.5	2.17	2.48	2.87	-0.01	0.23	0.70	-0.16	-0.31	0.39	-0.44	0.56
e4.0	2.38	2.69	3.07	-0.01	0.19	0.69	-0.16	-0.31	0.38	-0.45	0.55
e4.5	2.57	2.89	3.26	-0.01	0.14	0.69	-0.15	-0.31	0.38	-0.45	0.55
e5.0	2.75	3.07	3.44	-0.01	0.07	0.69	-0.13	-0.32	0.38	-0.46	0.54
e5.5	2.92	3.24	3.62	-0.01	-0.01	0.70	-0.12	-0.32	0.38	-0.46	0.54
e6.0	3.08	3.41	3.79	-0.01	-0.09	0.71	-0.10	-0.33	0.38	-0.47	0.53
e6.5	3.23	3.57	3.95	-0.01	-0.18	0.72	-0.08	-0.34	0.38	-0.47	0.53
e7.0	3.38	3.72	4.11	-0.01	-0.28	0.73	-0.06	-0.35	0.38	-0.47	0.53

cells show SDs (or ratios)

Takeaways:

- The median is very roughly $.5 + \log_{10} k$, to within about .25. Much more closely, it's roughly $-.5 + 1.6\sqrt{\log_{10} k}$. So for example with $k = 100$, i.e. 100 embryos, we get $-.5+1.6\sqrt{2} \approx -.5+1.6\times1.4 \approx 1.75$, close to the real value of 1.73.
- The 80% range between the bottom 1/10th and the top 1/10th is very very roughly 1. A closer approximation is $1 - \log_{10}(\log_{10} k)/4$. So for $k = 100$ we get $1 - (\log_{10} 2)/4 \approx 1 - .3/4 \approx .93$, sorta close to the real value of .81. But probably best to just remember that the range is greater than 1 SD for very small k, 1 SD for smallish k, and about .7 for $k >> 100$.
- General note: $10^{1/10} \approx 1.26 \approx 5/4$, and $10^{3/10} \approx 2$.
- The split of the 80% interval between 1/10 up to median, vs. median up to 9/10, starts at 50/50 and gets more lopsided as k grows, going to almost 40/60 for k around 100, and then returning to an evenish split. So we can basically think of it as 50/50, or 45/55 to be a bit more precise. We'll write e.g. that for $k = 10$, the result is about $1.05 \pm .53$ for the 80% range, even though really it's sometimes more like a 45/55 split. A simple qualitative summary to remember is that the results of embryo selection skew so that there are a bit more very-high outcomes than very-low outcomes. (I think what's happening is that, at first, the spread of the gaussian-max dominates. That spread decreases with more samples. But then the embryo-SD mixture spread starts dominating, and that is symmetric and doesn't decrease.)

Converting to IQ points, and showing the half-widths of 80% intervals:

all estimates from the median of $\mathbb{E}\_\sigma\ [\mathcal{N}(0,\sigma)]k$, which gives the raw selection power
the ± ranges describe the 80% range of outcomes between the bottom 1/10 and top 1/10 outcome.
(the number given is half that range.)

number (k) of embryos sampled. eX means 10^X

	$\mathbb{E}\_\sigma\ [\mathcal{N}(0,\sigma)]k$	effective IQ SDs	IQ points
1	0.00 ±0.90	0.00 ±0.36	0 ±5
2	0.38 ±0.74	0.15 ±0.30	2 ±4
3	0.57 ±0.67	0.23 ±0.27	3 ±4
4	0.70 ±0.63	0.28 ±0.25	4 ±4
5	0.79 ±0.60	0.32 ±0.24	5 ±4
6	0.86 ±0.58	0.35 ±0.23	5 ±3
7	0.92 ±0.56	0.37 ±0.23	6 ±3
8	0.97 ±0.55	0.39 ±0.22	6 ±3
9	1.01 ±0.54	0.41 ±0.22	6 ±3
10	1.05 ±0.53	0.42 ±0.21	6 ±3
e1.5	1.42 ±0.45	0.57 ±0.18	9 ±3
e2.0	1.73 ±0.40	0.69 ±0.16	10 ±2
e2.5	2.01 ±0.38	0.80 ±0.15	12 ±2
e3.0	2.25 ±0.36	0.90 ±0.14	14 ±2
e3.5	2.48 ±0.35	0.99 ±0.14	15 ±2
e4.0	2.69 ±0.35	1.08 ±0.14	16 ±2
e4.5	2.89 ±0.34	1.15 ±0.14	17 ±2
e5.0	3.07 ±0.35	1.23 ±0.14	18 ±2
e5.5	3.24 ±0.35	1.30 ±0.14	19 ±2
e6.0	3.41 ±0.35	1.36 ±0.14	20 ±2
e6.5	3.57 ±0.36	1.43 ±0.14	21 ±2
e7.0	3.72 ±0.37	1.49 ±0.15	22 ±2

cells show SDs (or IQ points)
IQ points assumes a PGS with R of .4
effective IQ SDs = (SDs of raw selection power) × .4 (effective SDs / raw SD)
IQ points = (effective IQ SDs) × 15 (IQ points / effective SD)

Takeaways:

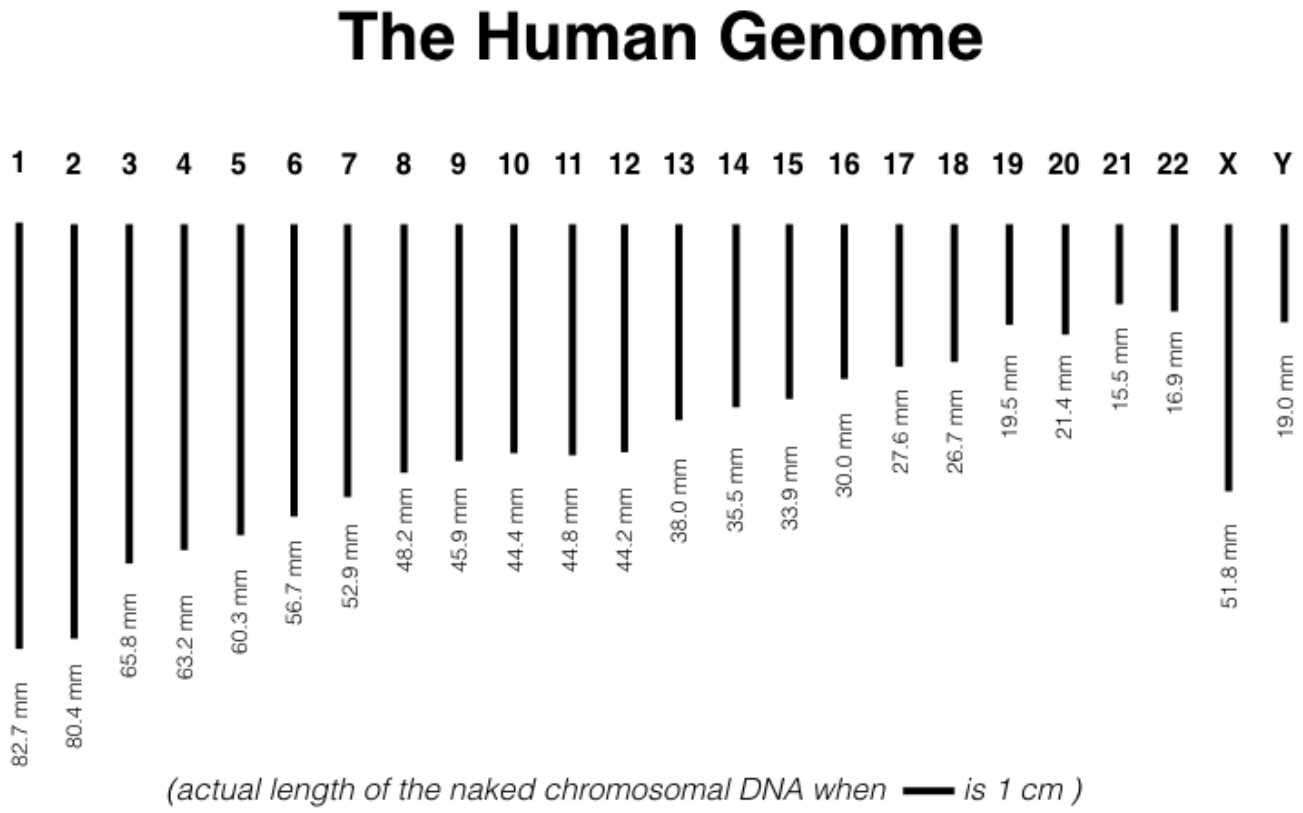
- The wall is real. With 10 embryos you get 6 IQ points. Up to 1000, each order of magnitude gets you +4 IQ points. Thereafter, each order of magnitude gets you +2 IQ points. So again: first OOM is 6 points, next two OOMs are 4 points each, further OOMs are 2 points each.
- Different couples will have outcomes that vary by about 3 IQ points up or down. (Again, all relative to the couple's mean.)

Appendix: Variation in chromosome length

(For this section, you can follow along in a python interpreter by importing this code, or just the lines at the beginning that define some lengths: https://gist.github.com/tsvibt/89cec1b4fd7d54be04ba34a0059bcc4c.)

The issue with chromosome lengths

Human chromosomes vary a lot in length:



(Figure 1.3 from Gallegos (2022)[171].)

This corresponds to a large variation in the number of base pairs in the chromosomes. The largest chromosome, chromosome 1, is over five times longer than the shortest, chromosome 21:

```
>>> max(chromosome_lengths.values())
247249719
>>> min(chromosome_lengths.values())
46944323
>>> max(chromosome_lengths.values()) / min(chromosome_lengths.values())
5.266871544829819
```

That means chromosome 1 has five times as much variance as chromosome 21. In several contexts in this article, we ignore this, and just pretend that all the chromosomes have the same length. For example, for simple chromosome selection of a haploid genome, we compute:

$$.5 \times 22/\sqrt{46} \approx 1.6$$

The .5 is a very rough estimate of the power, in SDs, of selecting the best 1 of 2 samples from a standard gaussian. We get to select 22 chromosomes (ignoring the sex chromosome). And then we pretend that each chromosome is 1/46 the length of the whole genome, so that it has 1/46 of the variance. That would mean a chromosome-SD is $\sqrt{1/46}$ of a whole-genome-SD.

Certainly in the limit, if one chromosome were really huge, like 99% of the genome, and all the others were tiny, we'd get very little selection power. The tiny chromosomes wouldn't matter, and the big chromosome would only give us about .5 or .6 SDs of selection power (divided by $\sqrt{2}$); it'd basically be gamete selection, were you only pick from 2 gametes.

Do chromosome lengths matter much?

What about for the actual chromosome lengths? Our rough formula, generalized, is this:

$$\text{(best 1 of 2)} \times \sum_{\text{chromosome } k} \text{chromosome } k \text{ SD}$$

[171] Gallegos, Maria. Fantastic Genes and Where to Find Them. Updated 2022-09-13. Accessed 16 February 2025. https://bookdown.org/maria_gallegos/where-are-genes/#preface.

The summation can be rewritten as:

$$\sum_{\text{chromosome } k} \sqrt{\text{Var}(\text{chromosome } k)}$$

Or:

$$\sum_{\text{chromosome } k} \sqrt{\text{chromosome } k\text{'s fraction of the genome}}$$

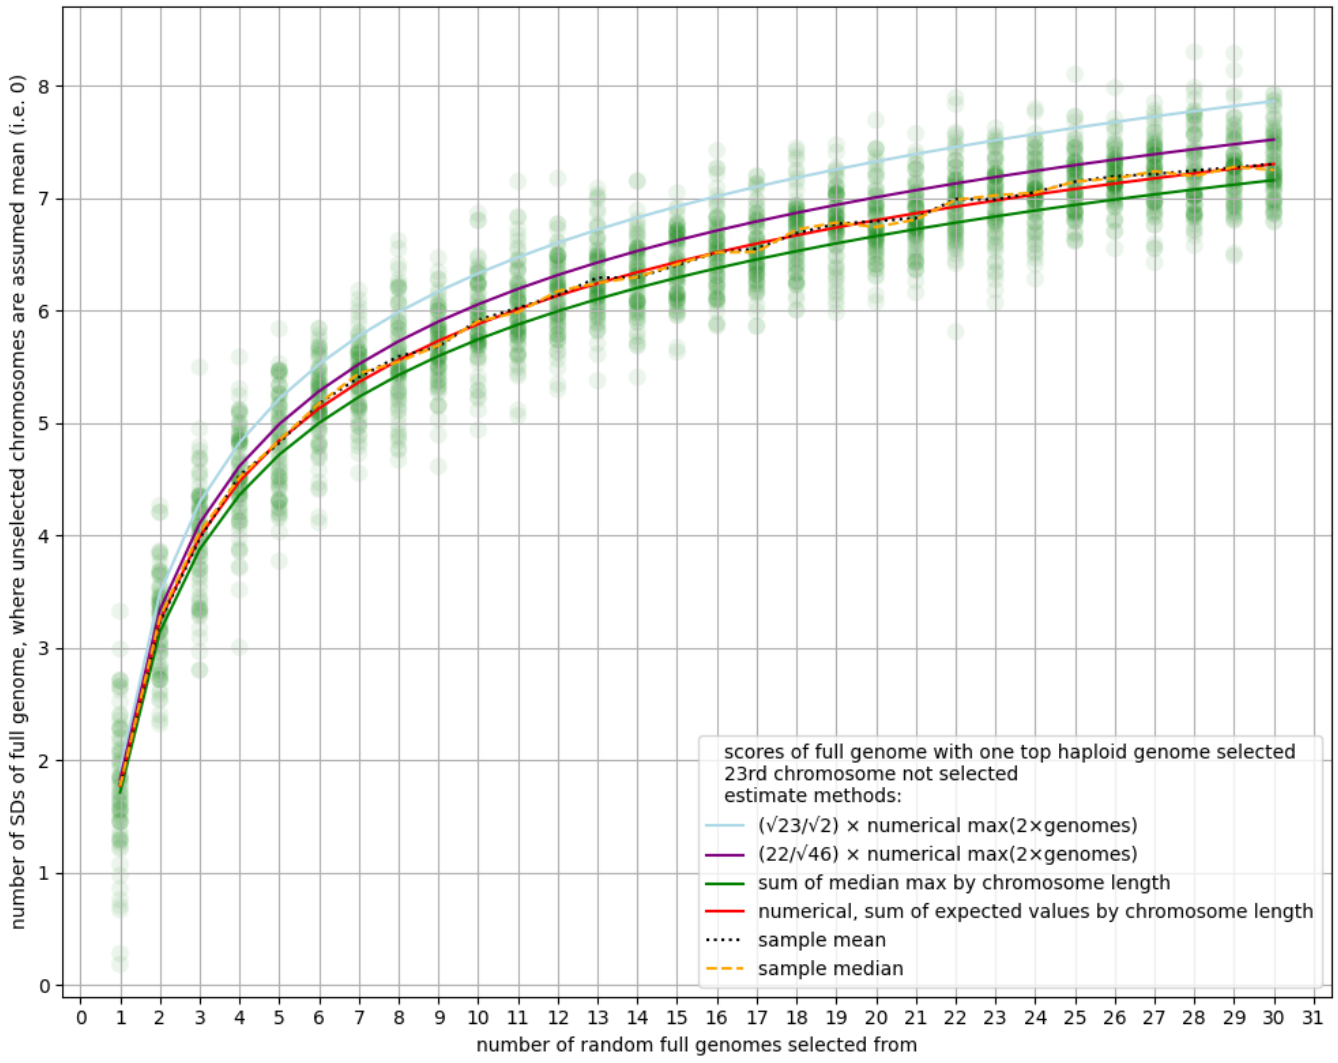
We compute:

```
>>> sum(np.sqrt(x) for x in [1/46]*23)
3.391164991562635
>>> sum(np.sqrt(x) for x in diploid_chromosome_fractions.values())
3.309001050655313
```

The first command gives the chromosome selection factor for 23 chromosomes of equal length. The second command gives the factor for 23 chromosomes of realistic length (where the 23rd chromosome has length the average of the lengths of the X and Y chromosomes).

The difference is less than 3%!

Here's a graph showing various ways of estimating chromosome selection, from "The power of selection":



Yes, you can see the difference between the estimates based on chromosome length and the ones that ignore it. You can also see how small the difference is.

Why don't chromosome lengths matter much?

The variance scales linearly with length, but the SDs scale as the square root of length. The longest chromosome is over five times as long as the shortest, but its SD scale is only a little more than two times the scale of the shortest:

```
>>> np.sqrt(max(chromosome_lengths.values()) / diploid_total_length )
0.20388105167623205
>>> np.sqrt(min(chromosome_lengths.values()) / diploid_total_length )
0.08883835371940446
```

To say it another way, the marginal returns on more SDs diminish with the square root, as you increase the length of the chromosome. Or to say it another another way, it's more important to select on the chromosome, than that it be long.

As a touchpoint: compare selecting just two chromosomes. They could be the same lengths, $(1/2, 1/2)$, or one could be 5 times longer, $(5/6, 1/6)$. For the former, the chromosome selection factor is $\sqrt{1/2} + \sqrt{1/2} = \sqrt{2} \approx 1.41$. For the latter, the factor is $\sqrt{5/6} + \sqrt{1/6} \approx 0.91 + 0.41 \approx 1.32$. Not even a 10% difference.

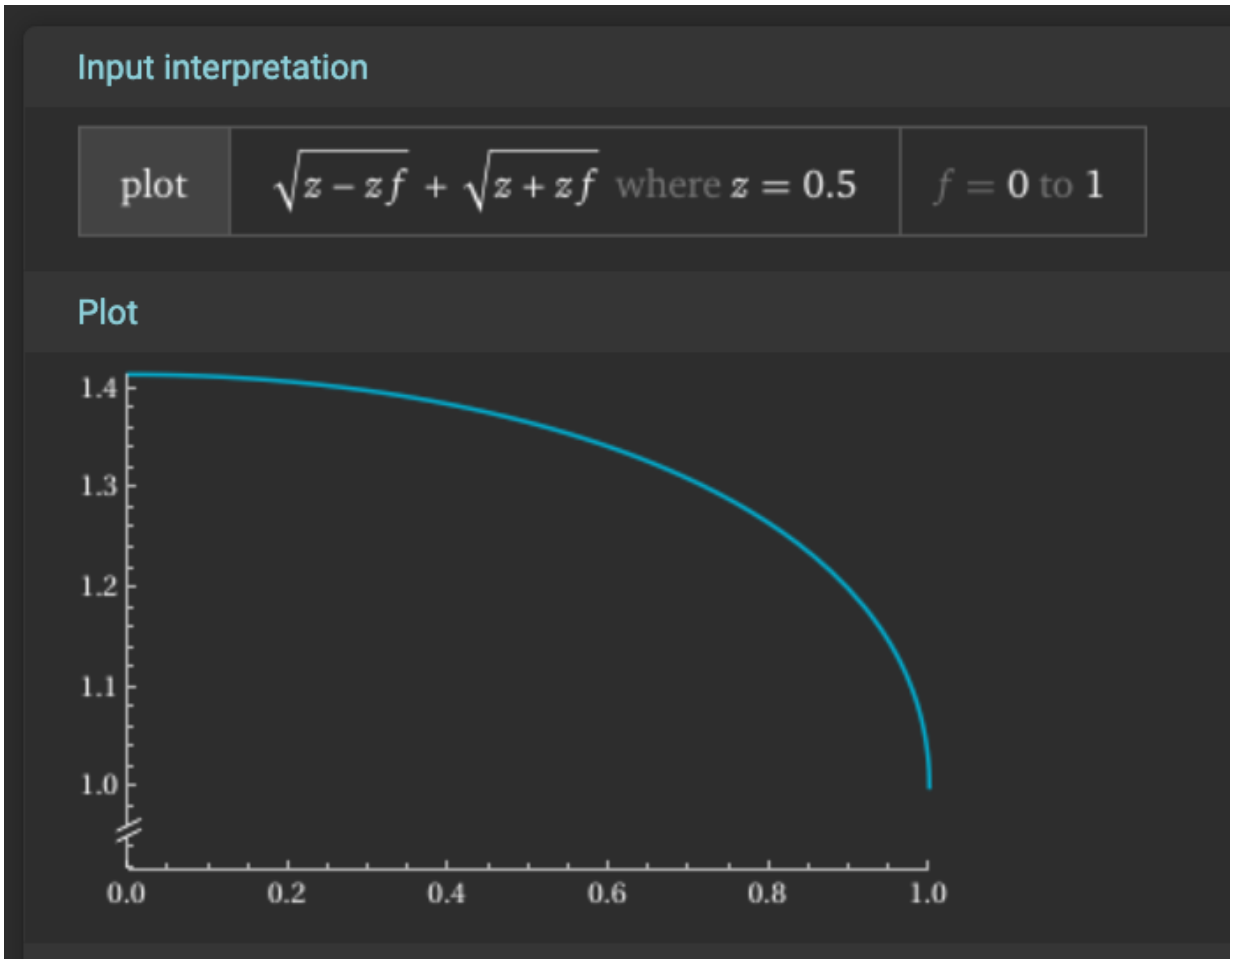
In fact, what about going all the way to $(1, 0)$? Then the factor is 1, which is about 70% of 1.41. In other words, completely merging two chromosomes of equal length is a 30% hit to their combined selection factor.

When *does* it make a big difference? Suppose we start with two chromosomes, each half the size of the whole genome so to speak, and then shift length from one to the other. Now we've removed some fraction f of one and added it to the other, so we have a $z + fz$ and a $z - fz$. As we move all the mass to the longer one, what's the selection factor?

We can plot this by giving the expression

```
plot sqrt(z - z f) + sqrt(z + z f) for z=.5, f from 0 to 1
```

to WolframAlpha. We get:



(The graph looks the same, but scaled, if you use $z = 1/46$.)

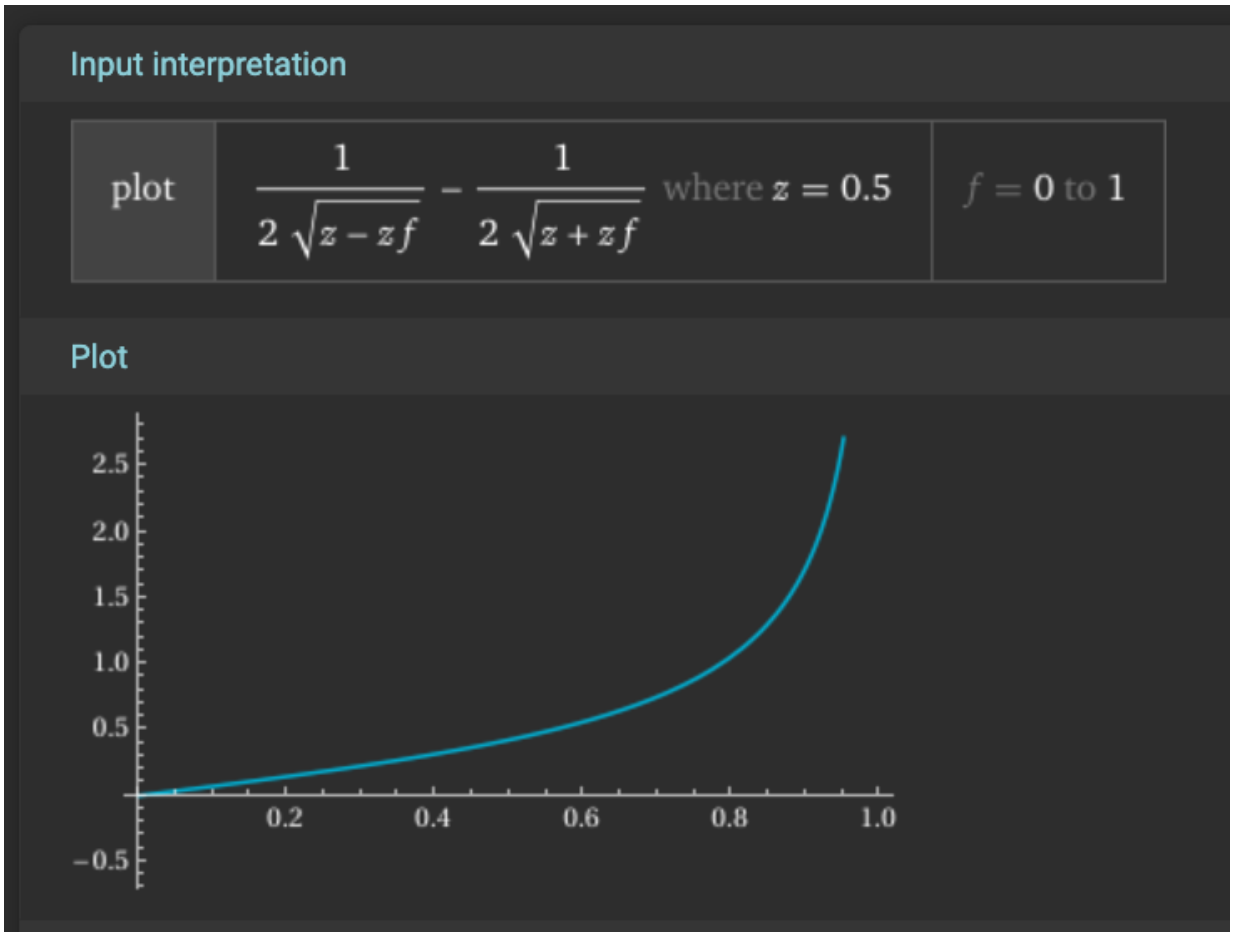
Another perspective: Say you're trying to decrease the power of selection by shoveling chromosomes from shorter chromosomes to longer ones, increase the length inequality. At first, starting from equal lengths,

you basically aren't doing anything: you're moving in opposite directions starting from the same point on a smooth curve, so at first the derivatives are just each other's negations. As the inequality increases, the difference in derivatives increases. But because of the shape of the square root, the losses from shortening the short guy only get *much* worse than the gains from lengthening the long guy when the short guy is pretty comparatively short.

To formalize this a bit, we can look at the derivative $\frac{d}{dx}\sqrt{x} = \frac{1}{2\sqrt{x}}$. Again suppose we start with both chromosomes of length .5, and then shift length from one to the other. If we keep doing that, what's the derivative? We plot this by giving the expression

```
plot (1/(2 sqrt(z - z f)) - 1/(2 sqrt(z + z f))) for z=.5, f from 0 to 1
```

to WolframAlpha. We get:



The rate of loss of SDs, as we shift length, starts at zero and climbs. But it only really spikes well toward the end, when we've almost eliminated the short chromosome.

Ignoring the 23rd chromosome

We did the above computations with 23 chromosomes. What if we ignore the 23rd chromosome? We get:

```
>>> sum(np.sqrt(x) for i,x in enumerate([1/46]*23) if i<22)
3.2437230354077378
>>> sum(np.sqrt(x) for i,x in enumerate(diploid_chromosome_fractions.values()) if i<22)
3.175290973758385
```

So with 23 chromosomes, the equal-length estimate is 3.39 vs. varying-length says 3.31. With 22 chromosomes, the equal-length estimate is 3.24 vs. varying-length says 3.17. I'd just as happily ignore the 23rd chromosome simply out of convenience to avoid complications (e.g., recombination works differently for the 23rd chromosome). Throwing out the 23rd chromosome from being selected over has more effect than taking lengths into account, and it's still less than 5%!

The upshot is that, for most selection purposes, it scarcely matters that the chromosomes are different lengths. Certainly if translated to real-world applications, none of our estimates are accurate to within 5% anyway.

Acknowledgements

Thanks to Sam Eisenstat for discussions of some of the mathematical ideas in this article. Thanks to kman for discussion of GV methods.

License

All content (such as diagrams) that isn't from an external source is shareable under the CC BY-SA 4.0 license (https://creativecommons.org/licenses/by-sa/4.0/) with attribution to the Berkeley Genomics Project, 2025.