[MUSIC PLAYING]

JOHN F. TISDALE: So I'm going to take the next 45-50 minutes to describe to you our efforts to bring curative therapies to sickle cell disease. I'll try to give you the broad strokes. If you have any questions along the way, please feel free to interrupt, and I can try to explain anything that didn't come off clear.

Lets start with the objectives. I'm going to get on my soapbox for just a second and say why think clinical trial participation is so important for developing curative therapies, and therapies in general for sickle cell disease. Then I'll describe what the current landscape is for bone marrow transplantation strategies in sickle cell disease. And then I'll touch on for a moment, if I have time, some of the novel curative approaches that we've been trying to develop based on what we know about bone marrow transplantation.

So I'd like to start with a description of sickle cell disease, and remind you that this is the first disease for which the molecular defect was identified by Linus Pauling and colleagues in 1949. And he went on to get the Nobel Prize for his work in protein biochemistry, and was later found a result from a single substitution at position six of the beta-globin chain, making an abnormal hemoglobin that's prone to polymerization upon deoxygenation. And this causes all the downstream consequences of sickle cell disease-- severe anemia, frequent vaso-occlusion with pain, and ultimately end organ damage and early death.

Now, current therapies are very limited. We only have one FDA approved drug to treat this disease-- hydroxyurea. Now, I think in large part due to the efforts by Mark at NIH, and subsequently here, there's been a growing interest in pursuing sickle cell diseases as a basic research disease or clinical research disease. And now we have a lot of growing interest in clinical trials in sickle cell disease.

So maybe I'm preaching to the choir here, but I'll just say a few more words about why I think that's important. And I think this slide really describes why clinical trials are so important when you're trying to develop a curative therapy for disease. So this is a slide from the New England Journal a few years back showing the dramatic improvement in overall survival in children with acute lymphoblastic leukemia who were enrolled in clinical trials. And you can see, in the late 1960s, cure rates were around 10%.

Over the ensuing decades, with clinical trials that started off with 400 or 500 individuals up to 8,000 individuals, those cure rates rose dramatically. So now a kid with ALL has a 95% chance of being cured of their disease. And I would argue that the urgency for treatment of childhood cancer-- these kids will die without treatment-- really encourages participation in clinical trials.

And I think we also need to develop curative strategies for sickle cell disease. This is Kathy Hassell's analysis of survival once we knew the denominator-- how many people are born in the US with sickle cell disease. Prior estimates of survival didn't include the complete denominator. This is after newborn screening was in place, so you can know how many kids were born with sickle cell disease and look at survival.

And you can see here that, in the 70s, survival was really quite poor in this disease in the US, with 10-15% of kids between the age of 1 and 4 dying early of overwhelming infections with encapsulated organisms, such that about 25% of kids were dying before reaching adulthood. This has really improved with the advent of newborn screening, and penicillin prophylaxis, and pneumococcal vaccination. But you can see here, even in the 2000s, early death is common, with many patients dying in their 20s and 30s. And that's virtually unchanged since the 70s.

So we've done a lot better in the pediatric population, and we haven't really done much in the adult population. And in this series, in 2006, the mean age at death was 39, and that's about where it stands today-- only 35% of patients alive beyond 45 years of age. So this is a really bad disease, with a lot of suffering along the way, and early death.

Well, it turns out there is a cure for sickle cell disease, and we've known about it for a long time. It's a bone marrow transplant. And this is really what we're trying to do with the bone marrow transplant here.

So here we have sickling that results from red cells that grow from progenitors and bone marrow stem cells that have this misspelled beta globin gene. So we've really sought to develop strategies to cure the disease based on either replacing this bone marrow stem cell-- getting it from somebody else who doesn't have the disease-- or repairing the genetic mutation in these bone marrow stem cells so that these red blood cells no longer sickle. So I'll show you what our strategy has been over the last couple of decades, and then I'll delve more into the individual strategies.

So first is allogeneic haematopoietic stem cell transplantation. So that just means you're getting a bone marrow transplant from somebody else-- in this case, we usually use a matched sibling donor. The patient undergoes some form of conditioning, we give them the bone marrow cells by vein, and get reconstitution with a normal amount of poiesis and normal red cells of donor type.

We've also been working on autologous haematopoietic stem cell and gene modification. So in this case, we take the patient's own bone marrow cells, put them in a flask where we expose them to a viral vector or an endonuclease that's attempting to either add a copy of the correctly spelled gene or edit the misspelled gene in those haematopoietic stem cells, and then give them back, like an autotransplant.

And then finally, there's been a lot of interest since the introduction of induced pluripotent stem cells as to whether or not we could make embryonic-like stem cells. Embryonic stem cells are derived from the inner cell mass of the blastocyst, and are encumbered by ethical problems in terms of moving them towards the clinic. But the ability to make embryonic stem cell-like cells from virtually any cell in the body is now available to us.

And so if we could take, for example, a skin cell, as shown here, reprogram it to an embryonic stem cell-like cell, correct the genetic mutations, make a haematopoietic stem cells, then we could do a transplant using the patient's own cells and fix this disorder. This is a bit further down the road, so I won't mention anything more about using IPS cells, except to say that it's on the horizon. And the first example in the mouse model where this was shown to be a potential curative strategy was in sickle cell mice.

So allogeneic transplantation is curative using a brother or sister when they're a match. And in fact, allogeneic transplantation in sickle cell disease was first performed in 1984, in a kid who not only had sickle cell disease, but had refractory acute myeloid leukemia. She was treated at St. Jude.

This is her as an adult. Now she's a spokesperson for transplant in sickle cell disease. She actually has her picture on the wall at St. Jude too. And she was cured both of her acute myelogenous leukemia and sickle cell disease with a standard bone marrow transplant using ablative doses of busulfan and siklofosfamid-- the standard way it's done in kids.

So Mark Walters and colleagues took this approach up in a multi-center trial, and reported, in the 90s, the first demonstration in a cohort of pediatric patients at risk for or having already had a stroke using myeloablative conditioning-- high doses of busulfan, high doses of siklofosfamid. They had an overall survival of 91% in an event-free survival. So a cure rate of 73% in this first series.

These results improved just with time. So just as you saw the cure rates going up in ALL over time, the same thing has happened in sickle cell disease. So this is the identical protocol applied to 87 kids in France, and their event-free survival improved to 95% after January of 2000. So just by doing it over and over, they got better and better at it, such that 95% of kids can now expect to be cured with an allogeneic bone marrow transplant from a matched sibling.

So that's very important, but the dilemma has always been-- kids do pretty well now with sickle cell disease, so most people don't want to do something as intense as a bone marrow transplant. But the flip side of that is, once they reach the adult clinic, it's a very different situation. There's accumulated end organ damage that precludes most adult patients from being able to get an allogeneic bone marrow transplantation. So this problem really had us all at a standstill. Kids are doing well, so they stopped doing transplants in kids for the most part. Adults are doing poorly, so they're not transplanting them because they're not eligible.

And then Mark Walters made an observation in a follow up of the original patients that suggested that you don't need to completely knock out the bone marrow to get a bone marrow to work in sickle cell disease. And this was based on long term observations in kids that were treated with busulfan and siklofosfamid when busulfan was a by mouth drug. It's very emetogenic, so very often kids would vomit when you're treating them, and you'd have to count the pills and put them back in hope that you got a full dose of busulfan. And In some of those cases, they didn't. The bone marrow wasn't ablated. And despite that, these kids were still cured.

And when they had a sex mismatch in the donor and recipient, and they could look at the blood to see how much of the blood was coming from donor, in one case, only 11% of the blood-- the white compartment-- was coming from donor. But despite that, all of the red cells were donor. So that suggested that, because the red cells of normals last so much longer than the red cells in sickle cell disease, this competitive advantage would allow you to do a transplant potentially with only getting 10% of the bone marrow replaced by the donor. So we wanted to see if we could intentionally achieve engraftment, without ablating the bone marrow. That would allow us to extend bone marrow transplants to sick adults with end organ damage who can make the decision for themselves.

So this is a cartoon of what we're trying to do. In the upper panels, you see a myeloablative transplant where the patient's cells-- in blue-- after conditioning, completely go away, and the donor cells-- in green-- completely replace, and then you would have reconstitution with normal red cells. In the non-myeloablative setting, you'd have some fraction of the recipient's cells that are ablated, leaving some behind, such that, when you reconstitute, you have a mixture-- in this case, 50/50 donor and recipient.

And we would expect, based on Mark's observations, that that would be enough to flip the red cells completely to donor. And so we asked the question, can we really do this on purpose, and that way, develop a less intensive conditioning regimen so that we can improve safety. So we did a lot of work in mouse and non-human primate, came up with a low-dose total body irradiation approach-- 300 centigrade, as shown here. And this is based mostly on Rainer Storb's work in Seattle showing that 200 centigrade cyclosporine can get a bone marrow graft into patients pretty robustly. But these were mostly cancer patients who had been previously treated with chemotherapy, so it's a little bit easier to get a graft in.

And Robbie Iannone, when he was at Hopkins, had started a very similar approach in individuals with sickle cell disease. So six subjects got the Rainer Storb's regimen-- 200 centigrade cyclosporine. They all engrafted-- had transient red cells coming from donor. They weaned immunosuppression with cyclosporine at day 100, and they all rejected.

So that told us that this wasn't going to be enough in sickle cell disease. We were going to have to do something further to get the grafts in. And we had been looking at rapamycin, now termed sirolimus as a way to make t-cells tolerant. And it's based on this oversimplification of the immune system here-- t-cells encounter antigenomic C complexes in the context of second signal, liberate cytokines that are required for effective function and proliferation.

So this is like signal one plus signal two equals immune response. Signal one, in the absence of second signal, renders t-cells anergic, and is required for the induction of tolerance. I show here cyclosporine blocks activation of the t-cell. So it's a really good immunosuppressive drug, but it also blocks the induction of tolerance.

And we had evidence for that in vitro, and we wanted to test this in in vivo and a mouse model to see if we could overcome this by introducing sirolimus, or rapamycin as I'm calling it here, which blocks IO2 dependent t-cell proliferation. So it really blocks the second signal, while leaving the first signal intact, at least theoretically, allowing for the induction of tolerance. And so we developed a rejection-prone mouse model and tested this approach, where we gave either cyclosporine or rapamycin-- sirolimus-- for 30 days, after a single dose of 300 rads of total body irradiation.

And you can see the cyclosporine-treated mice rejected their grafts, even before withdrawal of immunosuppression, whereas the rapamycin-treated mice reached stable chimerism levels of about 80% in their white cells. So this would be enough to fix sickle cell disease-- actually more than enough to fix sickle cell disease, if we could achieve this in humans.

And we did this in the sickle transgenic mouse model as well, and could show replacement by a donor type hemoglobin in this model. So we added sirolimus to the regimen-- so low-dose TBI sirolimus-- mimicking what we'd done in the mouse. But the one thing we really wanted to avoid in humans getting a transplant for sickle cell disease was to not switch one disease that's really bad for another disease that's really bad. And that's graft-versus-host disease

So graft-versus-host disease is a common complication in adult patients undergoing haematopoietic stem cell transplantation. In our center, at least half getting bone marrow transplants in the malignancy setting experience some form of graft-versus-host disease, and we really wanted to eliminate that as a possibility. And so we added alemtuzumab, an anit-lymphocyte antibody, to the preconditioning regimen because it depletes those cells of the recipient that might cause rejection and of the donor that might cause graft-versus-host disease.

And there had been clinical trials that showed that it was very effective at preventing graft-versus-host disease in the malignant transplant setting, but, unfortunately, relapse was more common. So this is not a problem that would be associated with a transplant in sickle cell disease, because we don't have a malignancy that we're trying to eradicate with the donor's immune system.

So this is the regimen. Relatively simple. Give an antibody for a few days, a single dose of irradiation, and then oral sirolimus, targeting levels of 10 to 15. And then we give an unmanipulated G-CSF-mobilized peripheral blood graft. We don't do anything to the haematopoietic product that we get from their donors.

So this shows the results in the first 10 subjects, accrued from 112 screened. That's an important point-- many of these patients came to us with a known donor in the families, so it's probably less than 10% of patients with sickle cell disease who actually have an HLA-matched family donor. Patients tolerated the conditioning regimen without serious adverse events.

It's a very different regimen than a standard or even a reduced intensity transplant regimen. 9 of the 10 successfully engrafted. They all had stable mixed chimerism, like in the mouse, so none of them went completely to donor This was associated with normalization of their hemoglobin. They had no sickle cell related events-- peri or post-transplant. And surprisingly, none of them developed acute or chronic graft-versus-host disease.

But at the time of this report, all of the patients remained on immunosuppression, and so we'd really like to do, as we had done in the mouse, determine whether they were indeed tolerant, and get them off immunosuppression to avoid any potential complications of the immunosupressive therapy. So we modified the protocol to allow weaning of sirolimus. We were aided by a few patients who, despite our best advice, didn't continue to take their sirolimus.

A couple in particular that we were having a really hard time getting their dosing right, when I'd finally ask, when's the last time you've taken your sirolimus. And the patient asked me, when was the last time I was here. Was that six months ago. So we realized that at least some patients could be fine without long term immunosuppression. So we just made the rule 50%. If you have at least half of your immune system-- CD3 cells-- coming from your donor, then we felt like it was more safe to wean your immunosuppression. Because if you flipped one way, we want you to flip to donor.

And then we evaluated 30 additional patients. And you can see here, from the table, that these are very sick adults with end organ damage, hepatopathy, cirrhosis, chronic bilirubin in the 20, end stage renal disease on dialysis. So these are sick adult patients undergoing this regimen.

And these are the results. In the expanded cohort, again, about 90% of the patients had a reversion of their disease. This shows their hemoglobin's going up over time. This improvement in hemoglobin was also associated with a reversal of the elevation of tricuspid regurgitant jet velocity.

So 10 of our patients had greater than 2.5, which Mark had identified as an indicator of pulmonary hypertension and, most importantly for us, an indicator of poor prognosis long term. So we actually used this 2.5 as one of the entry criteria for the protocol based on its prognostic value. And you can see that, after transplant normalization in the hemoglobin, the TR jet returns to normal.

This is a slide that I like to pause on, because there are a lot of misconceptions about patients with sickle cell disease being drug addicts. Many of them get treatment in the emergency room, which is not ideal, because the assumption is they're just coming in because they want narcotics And I can assure you that it's the last resort for the majority of these patients to go to the ER because most of them feel like they don't get treated in a way that's really conducive to helping them.

It's often they feel like they're being treated as a drug addict. Mark said this many years ago, observation in our patients that if you can break the pain cycle, you can get them off narcotics. And the problem is, in sickle cell disease, you know they have recurrent, severe, narcotic-requiring pain, so it's very hard to break that cycle. But if you can transplant them, and get rid of their sickle cell disease, and stop the vaso-occlusive crises, you can get them off pain medication.

So this is our IV morphine equivalent weekly, around 800 milligrams IV equivalent. And you can see we can taper that off to almost zero, in all of these patients post-transplant. The only patients who remain are patients that are riddled with severe osteonecrosis in multiple joints. And those patients have had to continue on narcotics, but at much lower doses.

It's also really gratifying to see that patients don't come back to the hospital anymore. So they're not coming back to the ER for pain medications anymore once you cure the disease. So this is a typical number of admissions for an adult patient with sickle cell disease-- 3 and 1/2. And this is in the year prior to transplant. And you can see that goes to 0.3 in the years after transplant.

So they're not coming back to the hospital. And some of these hospitalizations one for a kidney transplant in a patient who was on dialysis at the time, another for hip replacement for someone who wanted to get their hip replaced after the bone marrow transplant. So these are not admissions for pain.

We then wanted to know whether we could stop immunosuppression. So what happened about this rule of 50% chimerism? So on the left panel, you can see the percentage of cells coming from donor in the lymphocyte compartment-- so CD3 chimerism here.

And it's split between those who weaned in the dashed line and those who didn't wean in the solid lines. So you can see patients, once they reached 50% donor chimerism at one year, were eligible to wean their immunosuppression. And so those patients came off, and you can see their chimerism levels are stable here.

AUDIENCE: Sorry. Did all of the people who [INAUDIBLE] weaning at 50% successfully come off?

JOHN F. TISDALE: Yes.

AUDIENCE: And so none of them you had to restart because of the [INAUDIBLE]?

JOHN F. TISDALE: None of them. Right? And some of the ones that we're under also weaned on their own, but they're shown us the under. And you can see the CD14-15-- so the myeloid chimerism is indistinguishable in the two groups. So whether they're on or off immunosuppression, they're chimerism remains stable. The dip at the end is just because there's fewer patients that account for those time points.

So these results have now been replicated by another group using the exact same approach. Actually, They got our protocol and initiated it at their center there. And you can see 12 of the 13 patients had in engraftment with no mortality, no acute or chronic graft-versus-host disease, no serious extramedullary toxicity. They also measured quality of life, and demonstrated a pretty substantial improvement in quality of life. And at the time of their publication, four of their patients had also been successfully weaned from sirolimus.

So this experience, along with an experience in haploidentical transplantation-- which I don't really have time to cover today-- is being developed by Courtney Fitzhugh, who's a new tenure track investigator at NIH, allowed us to determine what level of chimerism do we really need to read. So Mark Walters was just looking at whole blood chimerism, so that doesn't really tell you about the myeloid compartment. We'd done studies looking at CD34 cells, or the bone marrow stem cells, versus myeloid cells, and we showed they match. So the myeloid cells tell us what's going at the stem cell level.

And we wanted to know what level of myeloid chimerism do we need to get to fix this disease, and we have a hematology fellow who used to be a quant guy. He's physicist doing mathematical modeling, and then he decided to go to med school and do a residency, and now he's a hematology fellow. We were rounding on the weekend, and he was asking, why is it that you only need 20%, or 30%, or whatever it is to make this disease go away?

And I had all along thought it was just red cell half-life, because a red cell half-life is so much shorter in sickle cell disease. It doesn't take much production to overcome that. But it could be that the patient cells are really beaten up with EPO and then not very sensitive to EPO anymore. So the donor cells would be more sensitive, and get an advantage that way. So he did a calculation in the back of someone's chart, and figured out that it's probably just half-life, But he did a more formal analysis after we looked at over 70 patients that we'd followed long term.

We'd had two patients that had gotten a half-match transplant and one that had gotten a sibling-match transplant that had a decline of their chimerism that was more prolonged. And we'd been measuring their chimerism every month for the first six months, and then every six months for the first couple of years, and then every year.

These three had a Fall that taught us exactly how much chimerism you need, because all three patients developed the rise in their hemoglobin S, and they had trait donors of greater than 50%, recurrent severe anemia, and recurrent sickle cell disease symptoms when their donor myeloid chimerism fell below 20%. We had one patient on the haploidentical transplant protocol-- the half-match protocol-- who's got 22% donor chimerism, and she's fine. So it looks like 20% is the mark that you need.

So Stefan made a mathematical model based on erythrocyte half-life. He used erythrocyte as a surrogate for a erythrocyte half-life-- the reticulocyte count. There had been some papers previously that showed that reticulocyte count correlates very well with red cell half-life.

And so these are the results in those patients. So here, you can see in the solid line, a fall of their percent donor chimerism over time, then a rise in their hemoglobin S over time, right at about the 20% mark. This is the predicted value from Stefan's modeling. And the other thing that taught us is that red cell half-life significantly affects the model.

So if we change the half-life from seven days, which is about what it is in sickle cell disease, to three days, it's a lot easier to fix this disease. So if you're trying something like gene therapy, you might want to first start in patients that have a really high reticulocyte count, because it's easier to overtake those cells. On the other hand, if patients have an alpha thal deletion or two, and they have an associated prolongation of the red cell half-life, it might be more difficult to fix them with either type of transplant.

So we've shown that moderate donor chimerism is enough to correct this disease, 20% seems to be the mark. Low toxicity allows us to apply this in really sick adult patients with severe end organ damage, cirrhosis, end stage renal disease on dialysis. This split or mixed chimerism in the absence of acute or chronic graft-versus-host disease says that we have achieved tolerance in these patients.

Of the 36 patients who were beyond a year and eligible for weaning, 27 are now off immunosuppression with no change in their chimerism. But most patients, as I mentioned at the start of this, don't have a donor in the family. It's like 10%. So we've done something good for the few, so we really need to figure out how to make this more widely available for the other 90% of patients with this horrible disease. So we searched unrelated donor and cord registries, and found them to be similarly limited. That's because of the genetics of African-Americans.

We were looking at haploidentical transplantation with post-graft siklofosfamid. I mentioned some of that. One day, you guys can have Courtney Fitzhugh back and give you an update on how that's going. She's just reopened that protocol now, so that's open for a accrual. Virtually everyone has a donor in that context, but it's a lot harder to make it work.

And we've been working for a couple of decades on gene therapy using the patient's own bone marrow stem cells. So how does that work, and where does that stand. So again, autologous hematopoietic stem cell gene therapy, as it has been called for many years before it was ever shown to be therapeutic-- we used to be more careful and call it gene transfer-- is accomplished by addition of a normal or therapeutic gene to the bone marrow stem cell of that patient.

So you use the patient's own bone marrow hematopoietic stem cells. You have to use a vector that integrates, so we use viruses from nature that naturally integrate into our genome. And in this way, since you're using the patient's own cells, there's no risk of graft-versus-host disease, and virtually no risk of rejecting. And so it's potentially available to all subjects.

So this is really how it works. It's These viral vectors are like the Trojan horse. We stuff these Trojan horses full of the beta-globin gene, correctly spelled, move those into the nucleus where you can unload this beta-globin gene and hopefully contribute to the production of a sufficient amount of hemoglobin A to out-compete the sickle.

And when Mark and I first started at the NIH, this was working really well in the mouse. I was very enthusiastic about curing sickle cell disease 20 years ago, because you could cure a mouse of virtually anything 20 years ago by taking their bone marrow, using a viral vector, putting in the correct gene, squirting it back in after a dose of irradiation. And it really looked like that was going to work in humans.

So we started doing that in vitro-- the same techniques with human cells. It worked really well. Just as well as it worked in the mouse. And then we started marking bone marrow stem cells in individuals who were getting a transplant for something else. So these were myeloma patients, mostly, getting a bone marrow transplant from myeloma.

We take half their marrow, put in a gene, then put it back with the transplant and see what the frequency of that gene was in their blood. And it was one in 100,000, dropped to one in a million, required very sensitive PCR assays to detect, and then went away completely. This is, of course, way too low to expect clinical benefit in any disorder, especially in sickle cell disease. So we really needed a better model for human haematopoiesis. The mouse had really let us down in that regard.

So we went back to develop the non-human primate competitive repopulation model because Art [INAUDIBLE], who was running the hematology branch at that time, had started gene transfer efforts, and had done a few transplants in the rhesus macaque, and the results were identical to what we saw in humans. So really high predicted efficiency. So 50% of the bone marrow stem cells, if you grew them out in colony assays or you grew them out in vitro, would have the gene that you put in. And then you do the transplant-- it's one in a million and it goes away.

So we ask a bunch of questions in this model. Were we losing these cells because we're putting a foreign gene? Back in the day, we were using the neomycin phosphotransferase, which is a bacterial antibiotic resistance gene. It's easy to use, easy to clone, but maybe, if you make a human cell express a bacterial protein that gets rejected. It turns out the non-human primate is tolerant to the neomycin phosphotransferase, if you put it in through a bone marrow transplant after ablation.

We had some difficulty with the early retroviral vectors based on the moloney leukemia virus. It had a tendency to integrate upstream, and transcription start sites turn on genes. And if it happened to land in front of an oncogene, it would turn on an oncogene. And so there were some early problems in this field with insertional mutagenesis which have been overcome by the current vector systems, but I can talk more about that as we go.

But the biggest problem that we had was that, we're taking these bone marrow stem cells out-- we're actually mobilizing them, as we do in humans-- putting them in a flask, throwing a bunch of growth factors in there, have it outside of the body for 96 hours, and then squirting it back in and expecting them to behave like hematopoietic stem cells. What mostly happened is that those cells would either die or differentiate. You can look in the flask and see neutrophils, and megacariocites, and platelets.

So most of the cells were differentiated. We spent a lot of time developing the culture methods to preserve hematopoietic stem cells during the transduction when we were trying to get the viral vector into the bone marrow cells. And the other big problem was all these viral vectors that we had developed-- if we put in the beta-globin gene with the locus control region elements that are required to direct erythroid specific expression of the beta-globin gene, which you need in a disease like sickle cell disease, it was too complex a payload for these viral vectors. They would just rearrange.

We never could get stable transfer using these early vectors. And this was solved by Michel Sadelain's group, who switched to HIV as a viral vector system for a completely different set of reasons. HIV infects non-dividing cells well. The early moloney leukemia viruses don't. They require cell division. And as I said, we were driving these cells to divide.

We'd make them differentiate-- they don't engraft. If we use a virus that doesn't require a division, maybe we can get in quicker, not drive with cytokines, get better engraftment. And it turns out it can do all that, and it can also package this more complex payload reliability, without recombining.

So here, Michel showed therapeutic hemoglobin synthesis and a beta thalassemia mouse model. So correction of beta thalassemia using an HIV-1 based lentiviral vector encoding beta-globin, with all the bells and whistles that drive erythroidic specific expression. And then Philip Labush showed correction of sickle cell disease very soon after that with an HIV-1 based vector in two different mouse models of sickle cell disease.

So now we at least have the expression and transfer problem solved with the HIV system. Again, the mouse doesn't predict that we're ready to go. We still have some other work to do.

And so we wanted to turn back to the non-human primates and do the preclinical testing of the lentivrial viral system, but then we had this really big problem. That is that old-world monkeys are not infected by HIV, and therefore, if you make a vector based on HIV, it's not going to work in our model that I was just bragging was so great at predicting what we ultimately saw in humans. So we had to figure out where this block is coming, try to overcome this block, and then develop the culture systems, the transduction methods to get high-level gene transfer in repopulating hematopoietic stem cells in this model.

After a lot of effort, we found it. The most significant restricting element for HIV transduction of hematopoietic stem cells in the monkey is TRIM5alpha. TRIM5alpha is like a cells own innate immune system. It targets HIV capsid when there's a species mismatch for degradation. So if you use HIV in the monkey, the HIV capsid gets targeted for degradation because of this mismatch. But there is a simian immunodeficiency virus which infects monkeys and efficiently, if made into a vector system, can transfer genes to their hematopoietic stem cells.

And this just shows how we make these vectors. So we have three different plasmids that we put into cells to make a virus-- the genome itself that contains the beta-globin gene and bits of the HIV capsid. And we always use a VSV-G envelope, because that envelope allows it to get into just about any cell.

So this is a typical HIV vector. This a typical SIV vector. And we thought, well, we can overcome this TRIM5alpha restriction if we just swap the SIV capsid. That allows us to continue to use all the HIV vectors that we've made, and just get around this restriction by swapping out the capsid for SIV.

And so this shows a typical competitive repopulation assay that we've developed in the rhesus to compare different conditions. These animal serve as their own control by virtue of the fact that we mobilize, collect their CD34 cells, split them equally between two conditions, and then infuse them back and see who wins. This is a standard HIV-1 based vector expressing YFP and chimeric HIV vector expressing GFP. We put equal amounts back in the animal after irradiation and see who wins.

And here, you can see the chimeric HIV vector is performing very well with 40% of neutrophils out 2000 days post transplant, as opposed to only a few percent coming from the HIV-1 based vector. And now, for the first time, we have red cells about 20% expressing this gene. And this is coming from only half the graft. So if the full graft had been fused then, theoretically, that would be more like 40%.

We compared these in successive pairs of animals, and then compared directly to SIV vector. And again, you can see equal contribution now from the SIV vector, which we know works well in the non-human primate, and the chimeric HIV. So this has allowed us to do comprehensive testing of beta-globin specific vectors.

And this is an animal that received the beta-globin driven GFP, so that gives us a chance to see whether we can derive specific expression of a gene in red cells in the animal. And these are about 20% of the animals red cells expressing green fluorescence protein. This is the protein that makes a jellyfish glow.

We've done some other studies to look at conditioning. How much do we have to condition, because we have to get ready to move this into the patients? Do we have to ablate when we do gene therapy, or can we not ablate, like we've done in the allo setting? And it turns out, at least for gene transfer, the more you condition, the better engraftment you get. So at least for the moment, we're leaning towards myeloablation in this setting.

So these are 10 gray-conditioned animals-- eight gray, six gray, four gray. You can see that much better results are achieved when we give an ablative dose of 10 gray. That's in terms of the percentage of GFP expression and the vector copy number. So this is a way that we can quantitate how many vectors do we have per cell. We have around 5 copies of the vector per cell at the ablative dose.

So we've moved through cell culture-- small animal, large animal-- so it's time to start phase 1 clinical trials in the human. And so we now have a protocol opened in sickle cell disease testing gene transfer. The only difference from what I've shown you in the animals is that we were actually collecting stem cells from the bone marrow, because the current method for mobilization with G-CSF is contraindicated in sickle cell disease because of serious and even fatal complications in the use of that drug in this disease. So we collect the bone marrow, and then we a lentivirus to transfer a normal copy of the beta-globin gene, get chemotherapy, busulfan, to get the patient ready, and then infuse and modify.

So this study is now actively accruing. It's open in France, and in the US. We're running the US side. So key enrollment criteria-- you have to be an adult on the US side, not so in France. And severe sickle cell disease by the conventional criteria-- multiple vaso-occlusive crises per year, chest syndrome, stroke, TR jet greater than 2.5. And this is a safety study-- safety and feasibility, but also there are efficacy end points.

And I'll show you some of the data so far. So this is at a median follow up of 11 months. This looks like a standard autotransplant. Patients get febrile neutropenia, stomatitis from busulfan, bacteremia crises, fever-- the sort of things you would expect doing an autologous transplant in subjects with sickle cell disease. But we haven't had any adverse events that were related to manufacturing of hematopoietic stem cells with the vector.

This shows vector copy number. I showed you in the animal that we had around five. This, unfortunately, is a bit lower in humans with sickle cell disease-- around 1-- and it drops after infusion. We see this also in the non-human primate-- a drop. This is probably because we're measuring progenitors here, and we're measuring stem cells here. But unfortunately, this drop is lower than what you would expect to be beneficial in every subject. We had one subject here with a vector copy number that's more respectable at about 0.25.

Here, you can see a connotation of the hemoglobin coming from vector. So it's spelled slightly differently than normal beta-globin, so you can pick it up by HPLC and quantitate it. So in green here, is vector driven hemoglobin production over time. And you can see these patients all have hemoglobin coming from vector at various levels approaching 20%.

There's one patient that makes it over the 20% mark because of fetal hemoglobin re-activation. And you can see those patients that are right around the 20% mark have the better hemoglobin-- they're around 11. So some evidence for efficacy as well in two of the seven subjects who've reached what we think is the threshold of 20%.

AUDIENCE: What was conditioning again?

JOHN F. TISDALE: Busulfan ablative.

AUDIENCE: Ablative?

JOHN F. TISDALE: Ablative. Yeah. We think we have to just give all the advantage to the infused cells.

AUDIENCE: So the radiation too?

JOHN F. TISDALE: No, just busulfan.

AUDIENCE: Going back to the previous slide-- is that mostly hemoglobin F? Is that what--

JOHN F. TISDALE: There are a couple of patients that have re-activated hemoglobin F. Yeah. So this patient in particular has a lot of hemoglobin F. This patient also-- a lot of hemoglobin F. And that's seen post-transplant. The younger you are, the better you get hemoglobin F re-activation.

AUDIENCE: Is that stable over time hemoglobin F?

JOHN F. TISDALE: It's been stable so far, but we only have about a year of follow up. And it's been described in the allogeneic setting, when it fails, fetal hemoglobin levels being high for about a year or two. So I expect that that's not going to be stable-- the fetal part.

AUDIENCE: Just a niave question for a non-transplanter. Why, if you're doing ablative, do you have any F's? Why do you recover after an ablative transplant?

JOHN F. TISDALE: Because the entire graft was not corrected.

AUDIENCE: Oh.

JOHN F. TISDALE: Right. So it's 50% transduction, and there's probably still some recovery, even after busulfan.

AUDIENCE: Go it.

JOHN F. TISDALE: This shows the patient in France. It was a young kid on chronic transfusion for a stroke. Here, you can see the total hemoglobin after the transplant. This is hemoglobin S coming back, either from graft or from leftover. This is the hemoglobin coming from vector-- the beta 87. So it's about half and half vector driven and endogenous hemoglobin S.

This patient did not have much of hemoglobin F response, and they have a hemoglobin total of 11.7 off transfusions. And so they look like they've been converted to a sickle cell trait. This is the first patient transplant.

And these are just some clinical parameters in that patient-- weaned off transfusions, they haven't been hospitalized for pain, and the reticulocytes in LDH are trending more towards the normal range now-- showing a near resolution of hemolysis. At least in the first patient, there's evidence of success, and in two others.

So this is feasible now. We can collect bone marrows. We can do lentiviral transduction. We can do transplants with the safety that you would expect from an autotransplant node gene therapy related event. We have production in all patients of this therapeutic hemoglobin, but we don't have enough to fix everyone. So we need higher levels of anti-cycling hemoglobin.

So we've been looking at a number of the steps in this process. I'll just go through a few, because it's getting close to the end of the hour. The harvest has been a difficult part. Bone marrow harvesting, and processing, and sickle cell disease is more problematic than what we thought.

We're actually moving now to plerixafor mobilization. There have been some studies that suggest safety of this drug for mobilization, as opposed to G-CSF, and we've done two patients so far with really robust mobilization.

AUDIENCE: Why don't you exchange transfuse and then G-CSF mobilization?

JOHN F. TISDALE: That's been done, and they still had crises. And we routinely exchange transfuse before we do any of this-- bone marrow harvest, plerixafor mobilization. But I think plerixafor is going to work. It's a very short half life, and you get peak mobilization at about six hours. So I think this is going to really help us. We've added--

AUDIENCE: [INAUDIBLE]

JOHN F. TISDALE: Right. You get mostly lymphocytes, yeah. It's much easier to process.

We've worked on improvement of vector copy numbers. So the in gray are the before the enhancers, PGE2, and other ways to get higher transduction efficiency in brown there. So you can see vector copy number on average is somewhere between 2 and 3. So I think this is going to be at a level that one might expect clinical benefit. This is a real product on the right, and that patient has just been infused with a vector copy number of 3.3.

We've done some other things to try and improve vector copy number. This is a little technical. I think I'll skip this, except to say that this is a typical vector for adding beta-globin to hematopoietic stem cells. So the HIV, LDRs, the locus control region, the beta-globin promoter, three prime untranslated region, and GFP-- just as a marker, you can stick beta-globin there.

But they are always backwards, because you need intron 2 to get good, level expression of beta-globin. So if you put it forward when you're making an RNA virus, the intron gets clipped out. As opposed to every other vector system out there, beta-globin is always backwards. And it's always a lower titer. It's always lower efficiency.

It's always a pain to make, because it's a hundredfold lower generally. With some improvements, we got up to about tenfold lower, but that means your postdoc is making a hundred times as many plates when they're making [INAUDIBLE] as any other vector.

So we just asked the question can we turn it around, and it turns out we can get tenfold higher vital titers by turning it around, and tenfold higher transduction efficiency by simply putting it back forward. But we lose intron 2 when we do that. So here, you can see that the titers going back up.

This is reversed. This is forward. This is reversed. This is forward. And This is a standard worker-bee vector of GFP.

And we get tenfold higher in the monkey. So this monkey got forward, green, reverse, yellow. So you can see, the percentage is tenfold higher. And this one got forward, yellow, reverse, green-- tenfold higher. And that's also true for vector copy number-- 10 to hundredfold higher, long term out two years in these animals. So these vectors work much better.

And we can get really, really good beta-globin expression in vitro. So this is a patient with sickle cell disease, peripheral blood mononuclear cells using our current in the clinic vector. So you can see this hemoglobin A here. And this is our optimized forward vector here. So pretty good expression of hemoglobin A using this forward orientation.

So I'm going to skip to the end, and skip the remarks about gene editing, except to say that the new tools that we have for cutting DNA exactly where we want to cut-- for example, with the CRISPR-Cas system-- has really improved the prospects for correcting the defect in hematopoietic stem cells. And now we're at about 20% correction of the misspelling, using these new tools.

AUDIENCE: John, perhaps you can show us that. Because we had Jennifer [INAUDIBLE] in here to present to the house staff, and we had a very positive session on gene editing. I'd be curious to hear what you're doing with that.

JOHN F. TISDALE: So to make gene correction work is a really tough prospect. So you've got to induce homologous recombination, which is extraordinarily low efficiency. You can get about one in a million, generally, if you try to do this.

So there have been several ways that people have thought about making this better. One, and the most important, is to introduce a site-specific DNA break, and then let the machinery turn on. Develop an efficient protein gene delivery system, and then optimize the template so that you get a really good result.

So what we did is we started with GFP-- GFP is so much easier. We said, OK, we're going to make this Cas9 system. We're going to tell it to cut GFP. And then we're going to make a delivery system and see if it works. So we made a GFP cell line, and then we looked to see how efficiently we made the lights go out. It's 80%-90%. It's really easy to do.

And then we said, OK, can we correct. So we said, can we make GFP into YFP. And so we delivered in that way. We could get about 40% correction of GFP to YFP-- proved it by sequencing. And then we said, OK, can we correct the sickle defect. And the way we did that is to put in GFP in our template, and then deliver Cas9 to cut at the misspelled spot. So if we get correction, the cells will turn green. And

The system that we developed was based on lentiviral vectors, because we know we can get engraftment with these vectors. We know how they behave. Everyone else is using electroporation, which kills hematopoietic stem cells. One of the reasons why we made vector's 25 years ago is because the other way to transfer genes was electroporation and that's really toxic for hematopoietic stem cells.

So we said, can we make a system to package this really cool stuff, and get it into the nucleus, and do this. So what we did is we encoded the Cas9 as a fusion protein with cyclophilin A. Cyclophilin A gets incorporated into the viral particle, and therefore we can hijack the cyclophilin A to get the Cas9 in for us.

We made it in both orientations, and tested it, and it works fine. And then we made a non-integrating lentiviral vector by deleting a part of the LTR so it doesn't get into the chromosome. And then we packaged that guide RNA with the Cas9 into this viral particle, and then we looked to see if we can get correction using various orientations and various ways of getting template in as well.

So here is what we get if we integrate everything. So it works really well if we integrate. We get 40% correction of the beta-globin mutation if we integrate the vector, but you don't want to leave your scissors out. It's like your mom always told you-- don't leave your scissors out, or you'll wind up cutting yourself. And you don't want to leave these scissors integrated in the genome where they're just ready to cut if they run into a guide.

So if we use a non-integrating template and an integrating Cas9, it works really well. But if we deliver the protein in this way that I told you-- hijacking the cyclophilin, getting it in as a fusion protein, we get almost 20% correction with a relatively easy to make delivery system. So we're now trying to work on ways to optimize getting more Cas9 in, because you can see the difference between integrating Cas9 and not integrating. Cas9 is really the limiting feature here-- we're not getting enough Cas9 in to do this efficiently enough to contemplate correcting the beta-globin mutation.

But the low-hanging fruit, at the moment, is that there are a number of genetic mutations that have been described that cause hereditary persistence of fetal hemoglobin. And so many, including ourselves, have looked to see if whether we just use these scissors to cut those genes, and see if we can get fetal hemoglobin back on. And that works very well.

So I think the first trials that you'll see coming from this kind of technology will be using the CRISPR Cas9, or zinc finger nucleases, or something else to cut, for example, the gamma globin promoter or the BCL11A enhancer, which have both been associated with repression of fetal hemoglobin. If you cut those and turn fetal back on, you can correct this disease. Nature has taught us that that's a strategy that works, because there are humans that have inherited these mutations and have inherited sickle cell disease who have no symptoms.

So we already know it will work if we can do it. So I think that's where the field will go next. So yes, we can cure sickle cell disease. Results have improved through successive clinical trials.

We need to do more clinical trials, with more of these approaches, to get closer to being able to deliver curative approaches to more patients with sickle cell disease-- not just the 10%. And patient participation is essential. We have the advantage of a very robust and well educated patient population that Mark followed at the NIH and accumulated over these years, and we've been able to partner with those patients to begin to test these approaches. So I'll stop there and acknowledge mostly the patients, families, and the caregivers that have had to look after them during these clinical trials, and take any more questions that you may have.

[APPLAUSE]

AUDIENCE: Two quick questions-- are you applying this to other hemoglobinapathy's as well, in any of these three techniques? And secondly, when you get the chimera count to 20%, can you reinduce and get them back up again?

JOHN F. TISDALE: So first question-- yes and no. So we've had some patients with thalassemia [INAUDIBLE] of these trials, and the results looked similar. But our referral base is really sickle cell disease, so it's been predominantly sickle cell disease.

The second question is whether we can boost. So those patients, that I told you, that drifted below 20% chimerism, they actually still had lymphoid chimerism. So we gave them a non-ablative busulfan dose, and another dose of Campath, and another unmanipulated dose of their donor, and their chimerism levels are now up and stable. So we can boost them. It turns out the boost is more intense than the original transplant, so it feels kind of funny to go call it a boost, but--

AUDIENCE: Great talk. Really great work. Two questions were about the gene modification. So when using the CRISPR Cas9, are you looking to just correct the single nucleotide, or are you looking for a larger correction? Because it's so much more efficient the smaller what you're asking it to do.

JOHN F. TISDALE: Yeah. So at first, we were doing larger, because we wanted to put in something to look. But we eventually want to get smaller and smaller, so that's part of the optimal template design that we're working to do. So our biggest bottleneck right now is getting enough Cas9 in the cells, and then we'll whittle away at the template, eventually, to improve efficiency.

AUDIENCE: And in terms for the engrafting-- have you considered [INAUDIBLE] antibodies, or some other things that might be-- [INAUDIBLE] 47-- that are more likely to be effective when there's not an immunologic barrier?

JOHN F. TISDALE: Yeah. So we're working now to try and move some of these antibodies into the non-human primates system. So what he's mentioning is that there are antibodies now that we can give instead of giving poisons that can knock out the bone marrow temporarily. One of the biggest benefits so far in the published literature and the presented work is that you don't have a dip in your counts when you give these antibodies.

So you give these antibodies-- you give the cells-- the platelets don't drop. The hemoglobin doesn't drop. The neutrophils don't drop. And then they switch completely over to donor. This hasn't been in a mouse so far. So there are now a couple of companies that have sprouted to bring these antibodies to clinical trials, and we've been working with Dave Skadden to potentially test this in the non-human primates. Yes?

AUDIENCE: Real quick question. You were talking about the dangers of giving G-CSF to patients with sickle cell disease of course, and I know sometimes, under extreme stress conditions, patients with sickle cell trait can also have sickling. For using high doses of G-CSF, have you ever had problems with the donors when they get the high dose G-CSF in terms of sickling, or complications, or anything?

JOHN F. TISDALE: We haven't. So before we initiated the trial, we did a controlled clinical trial comparing G-CSF mobilization and trait to age and race matched controls. We did symptom scores and everything, and mobilization was equivalent, if not better, in sickle cell trait, and symptom scores were indistinguishable between the two. And Now we have an experience of about 60 individuals that we've mobilized with G-CSF with sickle cell trait. Because half our donors have trait, and we haven't encountered any complications.

The NMDP still excludes sickle cell trait for mobilization, and we've been pushing them to get rid of that exclusion, but it still stands.

AUDIENCE: So just to share with the residents. I think you're really seeing history being made here. There is an arms race going on now. Pretty much every month, there's a nature science New England Journal paper, with gene editing or gene therapy, targeting hematopoietic stem cells.

And what's very unique about the sickle cell as a target is you can take the bone marrow cells out. You can do everything in vitro. You can then check for integration errors. You can sequence the genome to look for off-target effects, knocking out the wrong thing, and then, with a safer product, go into the patient with a transplant, which is more difficult to do for other corrections of systemic diseases.

So I think you're going to see, in thalassemia and sickle cell, the first successful gene editing. So it's a really exciting area, and John's obviously in the middle of it. So I--

AUDIENCE: So [INAUDIBLE] we're writing bits of other examples.

JOHN F. TISDALE: Yeah. CGD. Harry [INAUDIBLE] has had some success.

AUDIENCE: So anyway-- hopefully, he'll stay around, if anybody wants to come up to talk to him. Congratulations on this work, and I'm proud to have bumped shoulders with you in the ICU at the NIH over those years. So thanks.

[APPLAUSE]