Roots and other inverse functions

Injectivity, and introduction to inverse functions, especially square root functions

Remember that in a precise definition of the notion of a function f, the set E on which the function is defined is part of the definition. So for example, strictly speaking, the function f defined by $f(x) = x^2$ for $x \in \mathbb{R}$ and the function g defined by $g(x) = x^2$ for $x \geq 0$ are *different* functions. This distinction can be glossed over with impunity in much of elementary real calculus, but is essential for what we are about to consider, namely nth root functions in domains of the complex plane.

A key property functions may or may not have is called 'injective' or 'one-to-one': A function f defined on a set E is called injective, if it assigns distinct values $f(x)$, $f(y)$ to distinct inputs $x, y \in E$. In other words, if $f(x) = f(y)$ with $x, y \in E$ can only occur when $x = y$.
— Example: f given by $f(x) = x^2$ for $x \in \mathbb{R}$ is *not* injective, because for instance $f(7) = f(-7)$. In contrast, the function g given by $g(x) = x^2$ for $x \in [0, +\infty[$ is injective, because $x^2 = y^2$ for *nonnegative* x, y implies $x = y$. Similarly (re-using the letters f, g for new functions again), the function f given by $f(z) = z^2$ for $z \in \mathbb{C}$ is *not* injective, whereas the function g given by $g(z) = z^2$ for $\operatorname{Re} z > 0$ *is* injective.

In complex variables, when we restrict our attention to holomorphic functions on a domain, synonyms for 'injective' or 'one-to-one' have historically been used: univalent, or schlicht. Silverman's book uses the word 'univalent'.

When f is an injective function, defined on a set E with the set of values being called $f(E)$, the *image* of E under f, we can define an inverse function f^{-1} (defined on $f(E)$ with values in E) as follows: $f^{-1}(w) = z$ exactly if $f(z) = w$. If f fails to be injective, it does not have an inverse function, because $f(z) = w$ for a given w may not uniquely determine z, so there would be ambiguity in defining $f^{-1}(w)$.

Returning to the above examples, the function f given by $f(x) = x^2$ on $\mathbb{R}$ does not have an inverse function, but the function g given by $g(x) = x^2$ for $x \geq 0$ has an inverse function called the square root. Similarly the function f given by $f(z) = z^2$ for $z \in \mathbb{C}$ does not have an inverse function, whereas the function g given by $g(z) = z^2$ for $\operatorname{Re} z > 0$ does have an inverse function. g maps the right half plane $\operatorname{Re} z > 0$ on the set $\mathbb{C} \setminus]-\infty, 0]$, so its inverse function g^{-1} is defined for all $w \in \mathbb{C}$ that are not negative real numbers nor 0, and it has values in the right half plane $\operatorname{Re} z > 0$. There is a vast choice of possible domains G on which $h(z) = z^2$ would be injective, and a correspondingly vast choice of corresponding inverse functions h^{-1}. No single domain is a naturally preferred choice in all context, and this is why we are dealing with many square root functions in complex variables.

In real variables, we would in principle also have a vast number of choices (for instance $f(x) = x^2$ for $x \in [3.7, 9.35]$ is injective), but there are two natural choices with largest possible domains of injectivity: the square function on $[0, \infty[$ whose inverse function is $\sqrt{}$, and the square function on $]-\infty, 0]$, whose inverse function is $-\sqrt{}$. All purposes can be met by the single choice of the square function on $[0, +\infty[$, so there is no point in discussing multiple square root functions in real variable calculus, even though, strictly logically, this could have been done.

Definition: Assume n is an integer ≥ 2. We call f an nth root function on domain G, if $f(z)^n = z$ for all $z \in G$.

Outline

We have two main purposes here: To study root functions, and to understand the following general theorem:
Thm: *If f is holomorphic and injective on a domain G, then f' is automatically nonzero in G, and the image $f(G)$ (i.e., the set of all values $f(z)$ for $z \in G$) is itself a domain. We can then define the inverse function f^{-1} on $f(G)$, with values in G, and f^{-1} is automatically holomorphic. Moreover, the formula $(f^{-1})'(f(z)) = 1/f'(z)$ holds.*

One can begin with the general theorem and study root functions as a special case. This is the slicker approach. I will choose the somewhat clumsier path of doing root functions first, because it is a more hands-on and practical experience that requires less heavy abstract theory.

The first part of the theorem, that $f' \neq 0$ is typically complex-variable. In real variables, we have the example $f(x) = x^3$ for $x \in \mathbb{R}$. f is injective (and differentiable), but f' may vanish: $f'(0) = 0$. This is no contradiction to the complex case since $\mathbb{R}$ is not a domain (as it is not open in $\mathbb{C}$). If we were to fix this deficit by defining the cube function on an open set containing $\mathbb{R}$, we would forfeit injectivity, because a neighborhood of 0 will contain triplets of numbers ε, $\varepsilon e^{2\pi i/3}$ and $\varepsilon e^{4\pi i/3}$ that share the same cube ε^3. — Silverman's book postpones the proof to later, but we will use power series and an ad-hoc study of root functions to prove this part.

The second part of the theorem is technical: It guarantees that if f is holomorphic in a domain G, then $f(G)$ is again a domain, i.e., connected and open. Here the 'open' part is the difficult one, whereas 'connected' is easy: If a curve C in G connects z_1 with z_2, then $f(z_1)$ and $f(z_2)$ are connected by the image curve $f(C)$. The 'open' part is again typical complex variables. In real variables, $f(x) = x^2$ maps the open interval $]-1, 1[$ to the interval $[0, 1[$, which is *not* open. This claim that $f(G)$ is a domain again is known under the name 'Open Mapping Theorem' or 'Invariance of Domain Theorem'. Silverman's book has to postpone the proof for a few chapters; it is easier when we know $f' \neq 0$ already. Without knowing $f' \neq 0$ one can have a workaround using power series and root functions again.

Finally, the theorem stipulates that the inverse function is automatically differentiable. Again this is not true in real variable calculus, where the cube root function (as the inverse of the cube function) is not differentiable at 0, because the cube function has derivative 0 at 0. However, the complex cube function is not injective on any neighborhood of 0, so we are forced to remove 0 from the domain before we get an injective function that has an inverse. – The formula $(f^{-1})'(f(z)) = 1/f'(z)$ is of course well-known from real variable calculus already.

Complex root functions constructed in real terms

You know already how to solve $z^n = w$ for given w: Write w in polar coordinates $w = re^{i\phi}$, then the solutions z to the equation $z^n = w$ are given as $z_1 = r^{1/n} e^{i\phi/n}$, $z_2 = r^{1/n} e^{i(2\pi+\phi)/n}$, $\ldots$, $z_n = r^{1/n} e^{i(2(n-1)\pi+\phi)/n}$.

Each of these solutions, combined with the restriction to a suitable domain, gives rise to a complex nth root function, inverse function to an nth power function on a suitable domain. You will learn here, why there is no one-size-fits-all choice of a suitable domain that would allow us to select one definition as *the* nth root function; moreover, we have yet to see that our root functions are holomorphic, i.e., we have to check the Cauchy-Riemann differential equations. As mentioned, this is a bit clumsy, and a more abstract approach would avoid this detour through real variables. But I trust you will get a better practical grasp of complex root functions via the longer route.

Let's take as domain G the right half plane $\operatorname{Re} z > 0$ and write $z = x + iy$ as usual. We can convert z in polar coordinates $z = (x^2 + y^2)^{1/2} \exp[i \arctan \frac{y}{x}]$. We can define an nth root function f_0 on G by setting

$$f_0(z) := (x^2 + y^2)^{1/2n} \exp[i\tfrac{1}{n} \arctan \tfrac{y}{x}]$$

To check that f_0 is holomorphic in G, we note that it is real differentiable with (continuous) partial derivatives and confirm the CRDEs:

$$\frac{\partial}{\partial x}\, (x^2 + y^2)^{1/2n} \cos(\tfrac{1}{n} \arctan \tfrac{y}{x}) = \frac{\partial}{\partial y}\, (x^2 + y^2)^{1/2n} \sin(\tfrac{1}{n} \arctan \tfrac{y}{x})$$
$$\frac{\partial}{\partial y}\, (x^2 + y^2)^{1/2n} \cos(\tfrac{1}{n} \arctan \tfrac{y}{x}) = -\frac{\partial}{\partial x}\, (x^2 + y^2)^{1/2n} \sin(\tfrac{1}{n} \arctan \tfrac{y}{x})$$

I'll spare you the actual evaluation of these partials; it would be routine, albeit tedious.

We have constructed a holomorphic nth root function f_0 in the right half plane. For real positive z, this function reduces to the real root function $\sqrt[n]{z}$. Remember that the values of arctan are between $-\frac{\pi}{2}$ and $\frac{\pi}{2}$, so the root function f_0 takes values whose arguments are between $-\frac{\pi}{2n}$ and $\frac{\pi}{2n}$.

There is no natural reason why we should stay in the right half plane. However, the formula defining f_0 will cease to be useful when $x = 0$, because of the $\arctan \frac{y}{x}$ term. We could have written $\operatorname{arccot} \frac{x}{y} = \frac{\pi}{2} - \arctan \frac{x}{y}$ instead of $\arctan \frac{y}{x}$. This would have amounted to the same result in the first quadrant $x, y > 0$, and would have given us an nth root function

$$f_1(z) = (x^2 + y^2)^{1/2n} \exp[i\tfrac{1}{n}(\tfrac{\pi}{2} - \arctan \tfrac{x}{y})]$$

in the upper half plane $y > 0$. Recall that the arccot takes on values between 0 and π, so the values of g have arguments between 0 and $\frac{\pi}{n}$.

The same formula as for f_0 defines a holomorphic function f_* in the left half plane $x < 0$, but this function satifies $f_*(z)^n = -z$, not z, b/c arctan still takes values between $-\frac{\pi}{2}$ and $\frac{\pi}{2}$, whereas the argument of z in this plane is between $\frac{\pi}{2}$ and $\frac{3\pi}{2}$. So we have a root function in the left half plane given by

$$f_2(z) = (x^2 + y^2)^{1/2n} \exp[i\tfrac{1}{n}(\pi + \arctan \tfrac{y}{x})]$$

In the 2nd quadrant, f_2 coincides with f_1. We can continue this way and define

$$f_3(z) = (x^2 + y^2)^{1/2n} \exp[i\tfrac{1}{n}(\tfrac{3\pi}{2} - \arctan \tfrac{x}{y})]$$

f_3 coincides with f_2 in the third quadrant, but it does *not* coincide with f_0 back in the 4th quadrant. Rather, $f_4 = e^{2\pi i/n} f_0$ in the 4th quadrant.

The issue is geometrically easy to understand: as z moves around the origin (say on a circle), beginning on the real line with $\arg z = 0$ there, its argument increases and will have become 2π when z returns to the real axis. However, any nth root function we may choose, will have the argument of its value increase by a factor $\frac{1}{n}$ more slowly, starting (e.g.) at 0 as z is on the real line, and after coming around full circle, the value of that root function will have argument $\frac{2\pi}{n}$, inconsistent with the original value. This is why we cannot define a holomorphic root function on all of $\mathbb{C}$, or for that matter, $\mathbb{C} \setminus \{0\}$. We need to restrict the domain to some that disallows walking around 0.

Time to summarize the discussion in a picture (which is to scale for $n = 3$):

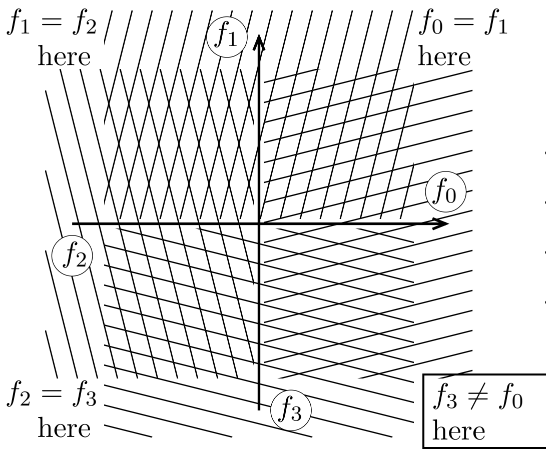

$$
\begin{aligned}
f_0(z) &= |z|^{1/n} \exp[i\tfrac{1}{n} \arctan \tfrac{y}{x}] &, \quad \operatorname{Re} z > 0 \\
f_1(z) &= |z|^{1/n} \exp[i\tfrac{1}{n}(\tfrac{\pi}{2} - \arctan \tfrac{x}{y})] &, \quad \operatorname{Im} z > 0 \\
f_2(z) &= |z|^{1/n} \exp[i\tfrac{1}{n}(\pi + \arctan \tfrac{y}{x})] &, \quad \operatorname{Re} z < 0 \\
f_3(z) &= |z|^{1/n} \exp[i\tfrac{1}{n}(\tfrac{3\pi}{2} - \arctan \tfrac{x}{y})] &, \quad \operatorname{Im} z < 0
\end{aligned}
$$

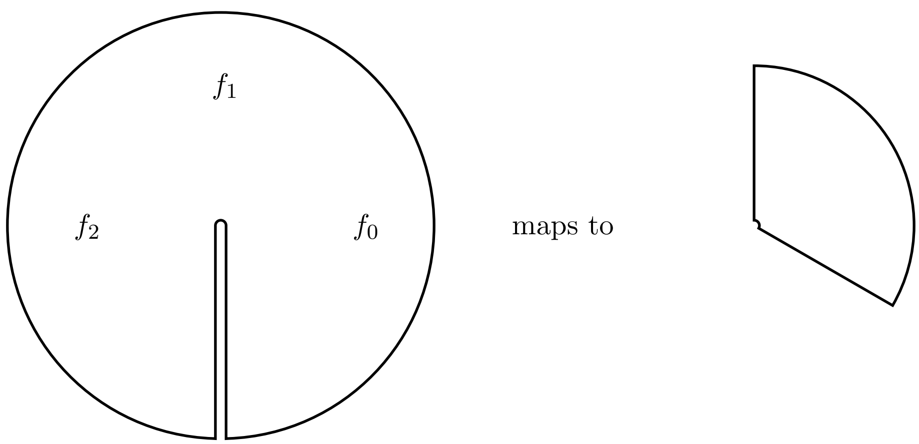

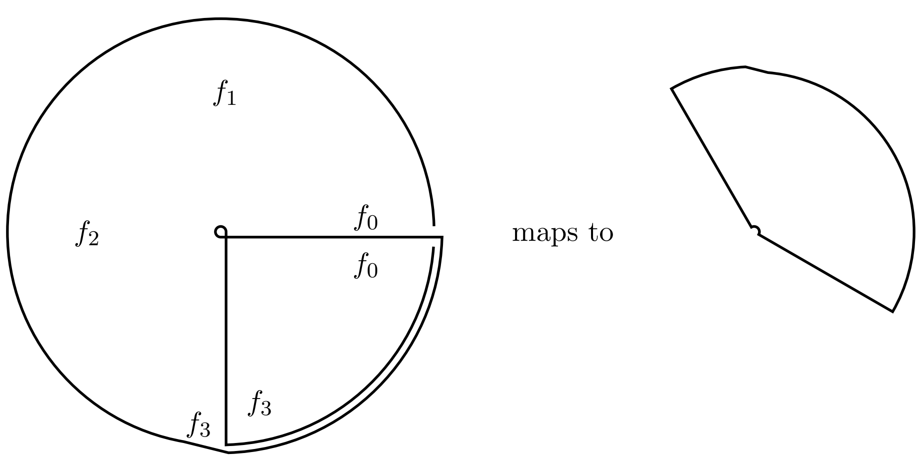

This is the opportunity to discuss which domains can actually carry an nth root function.

Fact 1: If f is one holomorphic nth root function in G, then $e^{2\pi ki/n}f$ for $k = 0, 1, \ldots, n-1$ are also holomorphic root functions in G, and there are no others. — This is clear since we know how to solve the equation $f(z)^n = z$ for $f(z)$; the only issue is that k doesn't depend on z; and this follows by continuity and the connectedness of G.

Fact 2: If $0 \in G$, then G cannot contain a holomorphic nth root function. — Proof: differentiate $f(z)^n = z$ and get $nf(z)^{n-1}f'(z) = 1$. Since $f(0)^n = 0$ implies $f(z) = 0$, the derivative at $z = 0$ would require $n \cdot 0^{n-1} \cdot f'(0) = 1$, an impossibility.

Fact 3: If G contains a circle going around 0, then G cannot contain a holomorphic root function. — We have seen this by deliberations about how the argument of z and a root function of z changes. The reasoning generalizes to arbitrary curves going around 0.

Fact 4: A simply connected domain that does not contain 0 does carry a holomorphic nth root function. — You'll see a prrof shortly, after we have done logarithm functions.

Logarithm ang general power functions

Def: We say f is a logarithm function in G if $e^{f(z)} = z$ for all $z \in G$.

[Only natural logarithms deserve to be considered. Stuff like $\log_{10} x$ is nothing but an abbreviation of $\ln x / \ln 10$ for the benefit of the uninitiated.]

We know how to solve $e^w = z$ for $z \neq 0$: If we write $w = u + iv$ then $e^w = e^u e^{iv}$; writing $z = re^{i\phi}$ in polar coordinates, we see that u has to be $\ln r$, and v has to be ϕ or differ from ϕ by a multiple of 2π. You have seen in a prior homework a logarithm function constructed in the right half plane by the formula $f_0(z) = \ln|z| + i \arctan \frac{y}{x}$; you have checked the CRDEs. The method of getting logarithm functions in the upper half plane, left half plane, etc., carries over, and again you get an inconsistency if you try to continue the process all the way around 0.

We get similar conclusions as for root functions:

Fact 1: If f is one holomorphic logarithm function in G, then $2\pi ik + f$ for $k \in \mathbb{Z}$ are also holomorphic logarithm functions in G, and there are no others.

Fact 2: If $0 \in G$, then G cannot contain a logarithm function (because e^z never vanishes).

Fact 3: If G contains a circle (or other Jordan curve) going around 0, then G cannot contain a holomorphic logarithm function.

These are proved exactly as in the case of root functions.

Fact 4: A simply connected domain that does not contain 0 does carry a holomorphic logarithm function.

Proof: Choose $z_0 \in G$ and some number w_0 such that $e^{w_0} = z_0$. (Such w_0 exist, because $z_0 \in G$, whereas $0 \notin G$). Define $f(z) := w_0 + \int_{z_0}^{z} \frac{d\zeta}{\zeta}$. The integral is independent of the path because G is simply connected and $\frac{1}{\zeta}$ is holomorpohic in G. It therefore defines a holomorphic function that we will call $L(z)$. We want to show that $e^{L(z)} = z$ in G. To this end we first show that $z^{-1}e^{L(z)}$ is constant:

$$\frac{d}{dz}\left(z^{-1}e^{L(z)}\right) = -z^{-2}e^{L(z)} + z^{-1}L'(z)e^{L(z)} = 0$$

since $L'(z) = \frac{1}{z}$. Now we calculate the constant by evaluating the expression at $z = z_0$: $z_0^{-1}e^{L(z_0)} = z_0^{-1}e^{w_0+0} = 1$. We have thus constructed L as a holomorphic logarithm function in G.

Once we have chosen a domain and a logarithm function on it, we can define the general power function on that same domain by $z^\alpha := \exp(\alpha \ln z)$. It is easy to see that for $\alpha = \frac{1}{n}$, this power function is an nth root function. If α is an integer, then every choice of logarithm function results in the same value for the power z^α, consistent with our old understanding of z^α as repeated multiplication.

In particular, with $\alpha = \frac{1}{n}$, we can use a holomorphic logarithm function in a simply connected domain G that doesn't contain 0 to construct a holomorphic root function there. This is an alternative approach to the one outlined before.

Policy: From now on, we consider usage of notations like $\sqrt{z}$, $\sqrt[n]{z}$, z^α, $\ln z$ for complex z as acceptable *if and only if* the context explicitly states or implicitly clarifies the domain of admissible z and a unique determination of the choice of root or logarithm function.

When possible, if the domain contains the positive real axis, we will deem it wise and convenient (albeit not strictly necessary) to choose root and logarithm functions that coincide with the real-variables definition on the positive real axis.

Example: 'Let $\sqrt{z}$ denote the square root function for $z \in \mathbb{C} \setminus]-\infty, 0]$ that has positive real values for $z > 0$.' In this context, $\sqrt{4}$ would equal 2, $\sqrt{\pm i}$ would be $(1 + \pm i)/\sqrt{2}$ resectively, and $\sqrt{-1}$ would not be defined.

Other example: 'Let $\sqrt{z}$ denote the square root function for $z \in \mathbb{C} \setminus i]-\infty, 0]$ (i.e., the complex plane minus the negative imaginary axis) that has positive real values for $z > 0$.' In this context, $\sqrt{4}$ would equal 2, $\sqrt{i}$ would be $(1 + i)/\sqrt{2}$, $\sqrt{-1}$ would be i, but $\sqrt{-i}$ would not be defined.

If $\ln z$ is a logarithm function in a domain G, then we obtain by differentiating $e^{\ln z} = z$ that $\frac{d}{dz} \ln z = \frac{1}{z}$.

If $\sqrt[n]{z}$ is an nth root function in a domain G, then by differentiating $(\sqrt[n]{z})^n = z$, we obtain $\frac{d}{dz} \sqrt[n]{z} = \frac{1}{n} \sqrt[n]{z}/z$.

Power series:

In this chapter, we take the nth root function in the right half plane that has real positive values on the real positve axis, and likewise we choose the logarithm function in the right half plane that takes real values on the positive real axis:

Then, using Taylor's formula and repeated differentiation at 1, we can obtain a power series representation

$$\ln(1 + z) = z - \frac{1}{2}z^2 + \frac{1}{3}z^3 - \frac{1}{4}z^4 + \frac{1}{5}z^5 - + \ldots = \sum_{n=1}^{\infty} (-1)^{n-1} \frac{1}{n} z^n \ , \quad |z| < 1$$

The radius of convergence 1 can be read off from the coefficients, but is also predicted by the fact that the nearest point to 0 where $\ln(1 + z)$ ceases to be definable as a holomorphic function is $z = -1$.

Given any complex number α, and k a nonnegative integer, we define the binomial coefficient

$$\binom{\alpha}{k} := \frac{\alpha(\alpha - 1) \ldots (\alpha - k + 1)}{k!}$$

There are k terms in the numerator. For $k = 0$ the above expression is to be understood as $\binom{\alpha}{0} = 1$ in accordance with the convention that empty products have the value 1. With this notation established, we get the binomial series:

$$(1+z)^\alpha = \sum_{n=0}^{\infty} \binom{\alpha}{n} z^n\ , \quad |z|<1 \ \ (\text{or for } 0 \le \alpha \in \mathbb{Z}\text{: } z \in \mathbb{C})$$

The radius of convergence 1 can be calculated from the coefficients with a bit of effort and Cauchy-Hadamard, but is more easily predicted by the fact that the nearest point to 0 where $\ln(1+z)$ ceases to be definable as a holomorphic function is $z=-1$; except, when α is a nonnegative integer. Then the power series terminates at $n=\alpha$ and $(1+z)^\alpha$ is a polynomial.

Interlude for purists:

Instead of appealing to the real variables functions, we could have constructed a root function in the domain $|z-1|<1$ directly by the binomial series:

$$z^\alpha := \sum_{j=0}^{\infty} \binom{\alpha}{j} (z-1)^j \ \ \text{for} \ \ |z-1|<1$$

To show that this series, for $\alpha = \frac{1}{n}$ is indeed an nth root function, we would need to prove first the formula

$$\sum_{k=0}^{j} \binom{\alpha}{k}\binom{\beta}{j-k} = \binom{\alpha+\beta}{j}$$

which guarantees that the Cauchy-Product for series z^α and z^β is indeed the series for $z^{\alpha+\beta}$. In particular, the nth power of the $z^{1/n}$ series is z.

This proof is not relevant for our purposes, so I just provide it for your reference, but am skipping it in class.

First we observe the relation $\binom{\alpha}{k} = \binom{\alpha-1}{k-1} + \binom{\alpha-1}{k}$. This is easily shown from the definition, by simplifying from the right hand side. With this lemma, we set up an induction proof: The claim is trivial for $j=0$ or $j=1$. Now for $j \ge 2$, we write an induction step:

$$\begin{aligned}
\sum_{k=0}^{j} \binom{\alpha}{k}\binom{\beta}{j-k} &= \sum_{k=0}^{j-1} \binom{\alpha}{k} \left\{ \binom{\beta-1}{j-k-1} + \binom{\beta-1}{j-k} \right\} + \binom{\alpha}{j}\binom{\beta}{0} \\
&= \sum_{k=0}^{j-1} \binom{\alpha}{k}\binom{\beta-1}{j-1-k} + \sum_{k=0}^{j} \binom{\alpha}{k}\binom{\beta-1}{j-k} \\
&= \binom{\alpha+\beta-1}{j-1} + \binom{\alpha+\beta-1}{j} = \binom{\alpha+\beta}{j}
\end{aligned}$$

Once this root function is constructed in $|z-1|<1$, one can extend the definition to larger sets by a variety of algebraic means, like for instance $(2^n z)^{1/n} = 2z^{1/n}$, which extends the definition from the circle $|z-1|<1$ to the larger circle $|z-2^n|<2^n$, and ultimately, by taking larger and larger n, to the entire right half plane. We won't bother to elaborate on this.

Inverse Functions in General:

Suppose f is holomorphic and injective in a domain G and f' does not vanish.[1] We claim that the inverse function f^{-1} is holomorphic on $f(G)$, in particular that $f(G)$ is a domain.

PROOF 1 (SKETCH), FOLLOWING SILVERMAN'S BOOK:

For those of you who are familiar with the implicit function theorem / inverse function theorem from advanced real multi-variable calculus, this is routine: The applicable variant of the theorem says that if we have a continuously differentiable real function f on an open set of $\mathbb{R}^n$ and a solution z_0 to the equation $f(z_0) = w_0$, and if the Jacobi matrix $DF(z_0)$ is invertible (i.e., its determinant non-zero), then there is a neighborhood V of w_0 and a neighborhood U of z_0 such that the equation $f(z) = w$ for $w \in V$ is still solvable, with a solution $z \in U$ that is unique there, and this solution $z = h(w)$ describes a continuously differentiable function h.

The proof of this theorem in advanced calculus basically relies on Newton's iteration method; and the fact that the Jacobi-Matrix is invertible makes Newton's iteration work. I omit technical details.

Now if f is holomorphic, the Jacobi matrix is $\begin{bmatrix} u_x & u_y \\ v_x & v_y \end{bmatrix} = \begin{bmatrix} u_x & u_y \\ -u_y & u_x \end{bmatrix}$ using the CRDEs. The determinant of this matrix is $u_x^2 + u_y^2 = |f'|^2 > 0$. So the inverse function theorem applies, and tells us that f^{-1} is *real* differentiable. The Jacobi matrix for f^{-1} is then the inverse of the Jacobi matrix for f and it can be seen from this that it again satisfies the CRDEs. Hence f^{-1} is holomorphic.

PROOF 2 (BY POWER SERIES):

I want to show you a proof of this theorem by means of power series. This proof levels the playing field between those that have and those that have not had senior level advanced calculus and stresses practical skills in working with power series, that are beneficial to both practically and theoretically-minded people:

Here is the overview: We first construct the coefficients for a power series representing f^{-1}, using a recursive algebraic calculation: **If** f^{-1} exists and is analytic, **then** its Taylor series can only be the one we are about to calculate. Next we show that the obtained power series has a positive radius of convergence and therefore defines an analytic function h in some (small) disk. Interpreted in this light, our previously formal calculation then means that h is indeed the inverse function f^{-1}. (The convergence proof is the tricky part).

In this argument we only use that $f'(z_0) \neq 0$ at a point z_0; the injectivity of f is not required; injectivity of f in a *neighborhood* of z_0 *follows* from this reasoning. Injectivity of f in all of G is a different matter and still needs to be assumed to get an inverse function on all of $f(G)$ rather than just on a small disk.

After this overview, let's go into details: f is injective on G, so f^{-1} exists and is defined on $f(G)$. Since f is holomorphic in G, we can write it as a power series in a neighborhood of an arbitrary $z_0 \in G$.

$$f(z) = \underbrace{f(z_0)}_{=:w_0} + a_1(z - z_0) - a_2(z - z_0)^2 - a_3(z - z_0)^3 - a_4(z - z_0)^4 - \dots$$

Here $a_0 = f(z_0)$ and $a_1 = f'(z_0) \neq 0$. For reasons of convenience that will become clear only

[1] As mentioned in the outline, it is an automatic consequence of 'holomorphic and injective in G' that f' doesn't vanish in G. But since we are not yet in a position to understand why this is the case, I am throwing the nonvanishing of f' in as an extra assumption here.

by later hindsight, I have chosen to call the coefficients beyond the linear term $-a_n$ rather than a_n.

Since this power series converges, there exists an R such that $|a_n| \leq |a_1|R^{n-1}$ (which is nontrivial only for $n \geq 2$): Indeed, for R_0 bigger than 1/(radius of convergence), we have $|a_n| \leq CR_0^n = (CR_0)R_0^{n-1}$ for some constant C. By making R yet larger, we have (assuming $n \geq 2$) the estimate $|a_n| \leq [CR_0(\frac{R_0}{R})^{n-1}]R^{n-1} \leq [CR_0^2/R]R^{n-1}$, and we can get $CR_0^2/R \leq |a_1|$.

So we have the following growth estimate for the a_n:

$$|a_1| =: A_1 \ , \quad |a_2| \leq A_1 R \ , \quad |a_3| \leq A_1 R^2 \ , \ldots \ , \quad |a_n| \leq A_1 R^{n-1}$$

We will now study the calculation of the coefficients of a power series $h(w)$, centered at $w_0 = f(z_0)$ with $h(w_0) = z_0$, that satisfies $f(h(w)) = w$. So we write

$$h(w) = \underbrace{h(w_0)}_{=z_0} + b_1(w - w_0) + b_2(w - w_0)^2 + b_3(w - w_0)^3 + b_4(w - w_0)^4 + \ldots$$

This time it is convenient to define the coefficients without the extra minus sign, as you'd expect anyways.

Now let's calculate $f(h(w))$ by plugging one power series into another. This will result in a glorious mess, that we can nevertheless handle skillfully (watch for the wisdom of indentation):

$$\begin{aligned} f(h(w)) = w_0 & \\ & + a_1 \left(b_1(w - w_0) + b_2(w - w_0)^2 + b_3(w - w_0)^3 + b_4(w - w_0)^4 + \ldots\right) \\ & \quad - a_2 \left(b_1(w - w_0) + b_2(w - w_0)^2 + b_3(w - w_0)^3 + \ldots\right)^2 \\ & \qquad - a_3 \left(b_1(w - w_0) + b_2(w - w_0)^2 + \ldots\right)^3 \\ & \qquad\quad - a_4 \left(b_1(w - w_0) + \ldots\right)^4 \\ & \qquad\qquad - \ldots \end{aligned}$$

The second line is a power series starting at power $(w - w_0)^1$; and I have written enough terms to eventually carry all powers up to order 4. The third line, once multiplied out, starts with power $(w - w_0)^2$; look carefully at it: I only need to carry terms up to b_3 if eventually I want to get all terms up to $(w - w_0)^4$. The fourth line, once multiplied out, begins with power $(w - w_0)^3$; and I only need to consider terms up to b_2 if I eventually want to get all terms up to power $(w - w_0)^4$. The fifth line begins with order $(w - w_0)^4$, and the (omitted) b_2 term will already not contribute to this order any more.

The upshot of this expansion is that the coefficient for each individual power $(w - w_0)^n$ can be calculated routinely (if tediously) by a finite algebraic calculation. This is because we centered the power series h that gets plugged into the other power series in such a way that no constant term remains.

We want to choose the b_n in such a way that

$$f(h(w)) = w = w_0 + 1 \cdot (w - w_0) + 0 \cdot (w - w_0)^2 + 0 \cdot (w - w_0)^3 + 0 \cdot (w - w_0)^4 + \ldots \ .$$

So we calculate the above mess, order by order and then compare coefficients:

$$\begin{aligned} f(h(w)) = w_0 &+ a_1 b_1 (w-w_0) \\ &+ \Big\{a_1b_2 - a_2b_1^2\Big\} (w-w_0)^2 \\ &+ \Big\{a_1b_3 - a_2\cdot 2b_1b_2 - a_3b_1^3\Big\} (w-w_0)^3 \\ &+ \Big\{a_1b_4 - a_2(2b_1b_2+b_2^2) - a_3\cdot 3b_1^2b_2 - a_4b_1^4\Big\} (w-w_0)^4 \\ &+ \dots \end{aligned}$$

The terms in braces must vanish, and a_1b_1 is desired to be 1. So we determine $b_1 = 1/a_1$ from the first line. Given the a's and b_1, we determine b_2 from the second line, then b_3 from the third line, and so on. In each step, to solve for b_n, we solve an equation $a_1b_n =$ *a combination of previously calculated terms*, since $a_1 \neq 0$, we can divide by it.

There is therefore a power series h that, *if convergent*, represents the inverse function $h = f^{-1}$, and we can practically calculate it by routine algebra (albeit tedious) to any order we like. However, we do not see a clear picture of a general formula for the nth term b_n emerging (and we can live with this ignorance).

We DO NOT WANT TO be more explicit about formulas for the b_n. That would be an awful mess, from which no insight can be gleaned! How then can we say that the b_n are such that the series h converges? This is the big miracle, and to force the good fortune, I had put the minus signs in front of the coefficients $a_2, a_3, a_4, \dots$. Because now, when I write the equation

$$b_n = \frac{\textit{a combination of previously calculated stuff}}{a_1}$$

the numerator on the right hand side is actually a polynomial with all *positive* coefficients, involving $a_1, \dots, a_n$ and the previously calculated $b_1, \dots, b_{n-1}$ (look at how the expressions in braces arise above, and you'll see it). If I plug the formulas for the previously obtained b_j in, these in turn would be polynomials with all *positive* coefficients, involving only the a_i and yet previously obtained b_i. In the end,

$$b_n = \frac{\text{poly}(a_1, \dots a_n)}{a_1^{N(n)}}$$

where $N(n)$ is a certain integer that we do not need to know in detail, and $\text{poly}(a_1, \dots, a_n)$ is a polynomial with all *positive* coefficients. And this is the info that we will live on.

Let's consider the specific example

$$F(z) = w_0 + A_1(z-z_0) - A_1R(z-z_0)^2 - A_1R^2(z-z_0)^3 - A_1R^3(z-z_0)^4 - \dots$$

For the solution series H to $F(H(w)) = w$ with $H(w_0) = z_0$, we call the coefficients B_n:

$$H(w) = z_0 + B_1(w-w_0) + B_2(w-w_0) + B_3(w-w_0) + B_4(w-w_0)^4 + \dots$$

The B_n arise from the A_n by exactly the same formula as the b_n arise from the a_n:

$$B_n = \frac{\text{poly}(A_1, \dots, A_n)}{A_1^{N(n)}} \qquad \text{where } A_n = A_1R^{n-1}$$

The following two features now save the day: (1) we can get explicit and simple formulas for the B_n even if we don't have nice formulas for the polynomial, namely $B_n = R^{n-1}/A_1^n$; and

(2) $|b_n| \le B_n$ for all n. With these two, we conclude the convergence of $\sum b_n(w-w_0)^n$ from the convergence of $B_n(w-w_0)^n$.

Namely for (2),

$$|b_n| = \frac{|\text{poly}(a_1,\dots,a_n)|}{|a_1|^{N(n)}} \le \frac{\text{poly}(|a_1|,\dots,|a_n|)}{A_1^{N(n)}} \le \frac{\text{poly}(A_1,\dots,A_n)}{A_1^{N(n)}} = B_n$$

where, at both $\le$ signs we have made use of the fact that the coefficients in the polynomial were *positive*.

And for (1): how do we know the B_n explicitly? Simply because we can explicitly sum the geometric series for F and then explicitly calculate its inverse function H, and then write H as a series:

$$F(z) = w_0 + 2A_1(z-z_0) - \frac{A_1(z-z_0)}{1-R(z-z_0)}$$

Solving $w = F(z)$ for z, we get a quadratic equation

$$2A_1R(z-z_0)^2 - (A_1 + R(w-w_0))(z-z_0) + (w-w_0) = 0$$

with the two solutions

$$z - z_0 = \frac{A_1 + R(w-w_0)}{4A_1R}\left(1 \pm \sqrt{1 - \frac{(w-w_0)8A_1R}{(A_1+R(w-w_0))^2}}\right)$$

where we use the square root function defined by the binomial series in a neighborhood of 1. We choose the minus sign in $\pm$ to get $z - z_0 = 0$ when $w - w_0 = 0$ and thus obtain a convergent power series $H(z)$ on the right hand side.

Now that we know that the series h converges, because the majorizing series H does, we know that h represents a holomorphic function in its disk of convergence; and this function is, by the same calculation, inverse to the function f.

Now let's suppose f is holomorphic and injective in a domain G and f' does not vanish there. Then the inverse function f^{-1} is defined on the set $f(G)$. Take any $w_0 = f(z_0) \in f(G)$. Writing f as a power series centered at z_0 (and convergent in some small disk about z_0), we can construct the inverse series h centered at w_0, and we know from the above that it converges. For w in some small disk about w_0, $z = h(w)$ is still sufficiently close to $h(w_0) = z_0$ and therefore in the domain G. So these w are in $f(G)$, and this shows that $f(G)$ is open. It is easy that $f(G)$ is also connected, hence $f(G)$ is a domain. And h is holomorphic in a neighborhood of any $w_0 \in f(G)$, hence is holomorphic in G.

This ends the proof.

Proof of the main theorem (from pg 2)

We repeat the theorem here:

Thm: *If f is holomorphic and injective on a domain G, then f' is automatically nonzero in G, and the image $f(G)$ (i.e., the set of all values $f(z)$ for $z \in G$) is itself a domain. We can then define the inverse function f^{-1} on $f(G)$, with values in G, and f^{-1} is automatically holomorphic. Moreover, the formula $(f^{-1})'(f(z)) = 1/f'(z)$ holds.*

We have to show that $f'(z) \ne 0$ in G. Then the theorem proved in the previous paragraph applies and gives us a holomorphic inverse function f^{-1} on the domain $f(G)$. The formula for the derivative of f^{-1} follows from differentiating $f^{-1}(f(z)) = z$ by the chain rule.

So to show that $f'(z) \neq 0$, we assume to the contrary that $f'(z_0) = 0$ for some z_0, and we then derive a contradiction from this assumption.

If $f'(z_0) = 0$, the Taylor series of f looks like

$$f(z) = f(z_0)+a_k(z-z_0)^k+a_{k+1}(z-z_0)^{k+1}+\ldots = f(z_0)+a_k(z-z_0)^k\left(1+\frac{a_{k+1}}{a_k}(z-z_0)+\ldots\right)$$

where a_k is the first *nonvanishing* Taylor coefficient beyond the constant term. Such an a_k must exist, because if all a_k were 0, then f would be constant and hence not injective. Note that $k > 1$ since we assumed $f'(z_0) = 0$. Choose some kth root b_k of a_k and some kth root function $\sqrt[k]{}$ in a neighborhood of 1 (say the one given by the binomial series. Then we can write

$$f(z) = f(z_0) + g(z)^k \quad \text{where} \quad g(z) = b_k(z-z_0)\sqrt[k]{\left(1+\frac{a_{k+1}}{a_k}(z-z_0)+\ldots\right)}$$

On some small disk D given by $|z - z_0| < \rho$, the function g has a holomorphic inverse g^{-1} by the power series construction of the previous section. The image $g(D)$ is a domain containing 0. It contains two distinct points $w_1 = g(z_1)$, $w_2 = g(z_2)$ (actually k of them), such that $w_1^k = w_2^k$; and $z_1 \neq z_2$ because g is injective. But now $f(z_1) = f(z_2)$ in contradiction to the injectivity of f.